IBM Internal Use Only

ZZ20-4167-7

Technical Bulletin

Performance Guidelines for IBM 3X74--Attached Workstations

by: Geert Bouman

IBM

IBM Internal Use Only

ZZ20-4167-7

Technical Bulletin

Performance Guidelines for IBM 3X74--Attached Workstations

by: Geert Bouman

Eighth Edition (December 1987)

This publication is an "IBM Internal Use Only" document, and therefore, its use is confined to IBM employees. The intent is to provide insight into the capability of IBM products for IBM marketing representatives, systems engineers, and industry specialists.

References in this material to IBM products, programs, and services do not imply that IBM intends to make these available in all countries in which IBM operates.

The performance data in this document apply to specific system configurations, transaction content, and prevailing system load conditions at time of measurement. The results which may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data in their specific environment. The data should not be used, either 'as is' or through extrapolation, to predict, guarantee, or in any other way to commit IBM to a specific set of performance results.

You may discuss pertinent information from this bulletin with a customer, and you can abstract pertinent information for presentation to your customers. However, the charts and tabular data are for internal information purposes only and may not be given to customers.

In this document, any references made to an IBM licensed program are not intended to state or imply that only IBM's licensed program may be used; any functionally equivalent program may be used instead.

Publications are not stocked at the address given below. Requests for IBM publications should be made to your IBM representative or to the IBM branch office serving your locality.

A form for reader's comments is provided at the back of this publication. If the form has been removed, comments may be addressed to V. J. Balachandran, IBM Corporation, Department G37, Building 002, Research Triangle Park, NC 27709. IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

PREFACE

Preface and Acknowledgments

This publication provides performance guidelines and data on the IBM 3174 and 3274 control units with a variety of workstations and devices attached. For a description of the characteristics and operation of these units, refer to the **IBM 3174 Subsystem Control Unit Functional Description,** GA23-0218, and the **3274 Control Unit Description and Programmers Guide,** GA23-0061.

This document is the result of the work of many people in the 3X74 Subsystem Performance Group in Kingston. Russell J. Houldin and So S. Chang developed and maintained all 3274 and 3174 performance models except for the Token-Ring Network connection. They were also responsible for measurement analysis. James T. Zahorsky and Jum L. Chin developed the performance measurement methodologies for all products, and obtained and analyzed many of the data. Tom F. Dubois and Richard A. Swanson were responsible for the models of the IBM 3174 Subsystem Control Unit Token-Ring Network connection and file transfer operations, and performed the measurements, analyses, and simulations.

Chapters 3 and 4 in this edition were expanded with information provided by Mr. Houldin, who also originated chapter 12 and the IBM 3193 data sheet.

I am deeply indebted to all of them for their contributions and reviews.

Geert H. Bouman Kingston

CHANGES AND ADDITIONS TO THE EIGHTH EDITION

The organization of this eighth edition has been modified to accommodate new information, and electronic publication of **IBM 3X74 performance data sheets** on HONE. (Performance charts cannot be distributed electronically at this time.) The purpose of these data sheets is to speed up the availability of new or improved performance data on a specific product to the field.

The seventh edition of this document, which was electronically distributed on HONE, appears as Chapter 10 in this edition.

Twelve chapters and appendixes introduce the performance aspects of various 3X74 subsystem operations, and illustrate how the data in 3X74 performance data sheets may be used. Refer to Appendix A for data sheets on 3X74-attached IBM products.

Curves and data in this edition **do not** include significant contingencies unless expressly indicated (for charts, by an asterisk). They have been obtained from models based on actual measurements. (When using information in the first through the fifth editions, be aware that most of that data includes 10-15 percent for contingencies.)

Furthermore, there is new performance information on the following products:

- The Asynchronous Emulation Adapter (AEA) of the IBM 3174 Subsystem Control Unit. (Amended performance information of IWSY Flash F8738, distributed on HONE.)
- IBM 3194 Display Station.

For the performance considerations and data on a specific device, use the applicable performance data sheet and consult the applicable chapter(s) and appendixes as needed.

Information on the following products has been updated:

- IBM Token-Ring Network 3270 gateway feature: performance of remote 3174 control units with gateway added
- Section on Multiple Logical Terminal operation in Chapter 2 has been extended for operation in CUT mode
- IBM 3270 PC and AT file transfer. Chapters 7 and 8 in the sixth edition have been combined in a performance data sheet for Personal Computer Workstations.

Some information in the sixth edition (ZZ20-4167-5) has not been included in this document, that is, on PC/VM Bond Release 2 (3274 data only), and on the IBM Personal Computer AT/370 (3274 data only).

For performance information on 3X74-attached graphic workstations, consult the **Graphic Systems Performance Guide,** ZZ20-5571, by A. J. Kirkland (IBM Hursley Laboratory). For recommendations on handling long data streams, as may be encountered in graphics applications, see Chapter 4 of these guidelines.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

CHAPTER 1. INTRODUCTION

This publication is intended for IBM personnel who need information on IBM 3174- or 3274-based subsystem performance. The data in these guidelines can be used for making performance estimates for 3X74 subsystem control units with devices of many different types attached. These not only include IBM 3270 display stations such as the 3191, 3278, 3278, and and 3290, but also IBM printers, and IBM Personal Computer (PC) and Personal System /2 (PS/2) products.

The purpose of these estimates is to help you select the most suitable IBM 3X74 subsystem configuration for your proposed or installed applications, when performance is a dominant consideration.

The performance of channel-attached subsystems and those with an IBM Token-Ring Network is dominated by the characteristics of the subsystem control unit. The performance of subsystems attached by a telecommunications link, on the other hand, is paced by the capacity of the link. For this reason, these 3X74 guidelines focus mostly on control units attached to a channel or Token-Ring Network.

With the advent of long data streams for graphics and image applications, file transfers, and high speed printing, the estimation of subsystem control unit utilization is important in determining subsystem configuration.

The data and charts in this document may be used for comparing the performance of selected configurations and operating modes. They can also be used to get a 'feel' for how the type and number of devices attached to a control unit will affect subsystem response time and its utilization. With the data on all other time delays in the system, a total response time estimate can be made. Control utilization is especially important in considering the effects of file transfer and high-speed printing on subsystem performance.

Most information presented in this document is based on actual measurements, and does **not** include significant contingencies unless explicitly indicated (for charts, by an asterisk). This should be kept in mind when making comparisons with information in editions one through five (ZZ20-4167-0 through -4), where usually a 10-15 percent contingency **was** included.

The number of digits used in the representation of data does not imply a degree of accuracy. Occasional discrepancies in the data may be due to slight environmental differences during time of measurement, for example, in acknowledgments. The performance data apply for microcode levels as indicated.

A more exhaustive performance analysis can be made using HONE AID FIVE3270 for the devices which are supported.

Some workstations that attach to an IBM 3X74 control unit can perform file transfer operations with the host. The effect of file transfer on the 3X74 is different from normal interactive traffic, and depends on the specific file transfer programs being used in the host and the workstation. Since there are several different file transfer programs, you should not use information on one environment to estimate the performance in another.

SUBSYSTEM RESPONSE TIME

In any interactive system, an inquiry or other message entered by an operator encounters multiple delays, as shown in Figure 1 on page 3, before a response becomes available on the display screen. The aggregate delay is commonly referred to as "response time", "system response", "user-perceived response time", etc.

These delays represent time expended in moving the inbound and outbound messages from place to place, time spent by various hardware and software components in executing the processing functions involved, and time spent waiting for various hardware and software resources to begin acting on a message.

The time spent in moving messages is proportional to their length and the speed of the transmission facilities. The time spent executing processing functions is related to the complexity of the work to be done. The time spent in waiting for a component is heavily influenced by its utilization, that is, the fraction of time it is busy performing work. As utilization increases, waiting time tends to increase progressively. For example, when subsystem utilization increases as a result of more terminals, higher transaction rates at the terminals, file transfers, and/or high speed printing, the utilization of the slower links in this chain, for example, a telecommunications line, will rise rapidly and slow response times appreciably.

In these guidelines, the subsystem **performance**, or Main Frame Interactive (MFI) **response time** of a display subsystem, is viewed as its contribution to the **user-perceived** response time experienced at an interactive workstation. User-perceived response is defined as the elapsed time from pressing the Enter key or a Program Function key until the **last character** of the response has been written on the screen, unless otherwise noted.

Thus, **subsystem response time** is the processing (and queueing) time in the subsystem of the inbound and outbound messages of a transaction. Delays in the host and the network (not addressed in these guidelines) are **not** included, unless specifically stated.

Subsystem performance, that is, the response time of a specific terminal/control unit combination, depends on:

- The subsystem control unit type -- its hardware and microcode
- The host-to-control unit link -- channel, telecommunication link speed, SNA or non-SNA
- Whether or not the subsystem includes a Token-Ring Network
- Workstation type and operational mode (Control Unit Terminal, or Distributed Function Terminal)
- The length and content of messages being processed
- The average utilization of a subsystem control unit as a result of MFI transactions, file transfers and printer operations, commensurate with the type and number of other devices in the cluster.

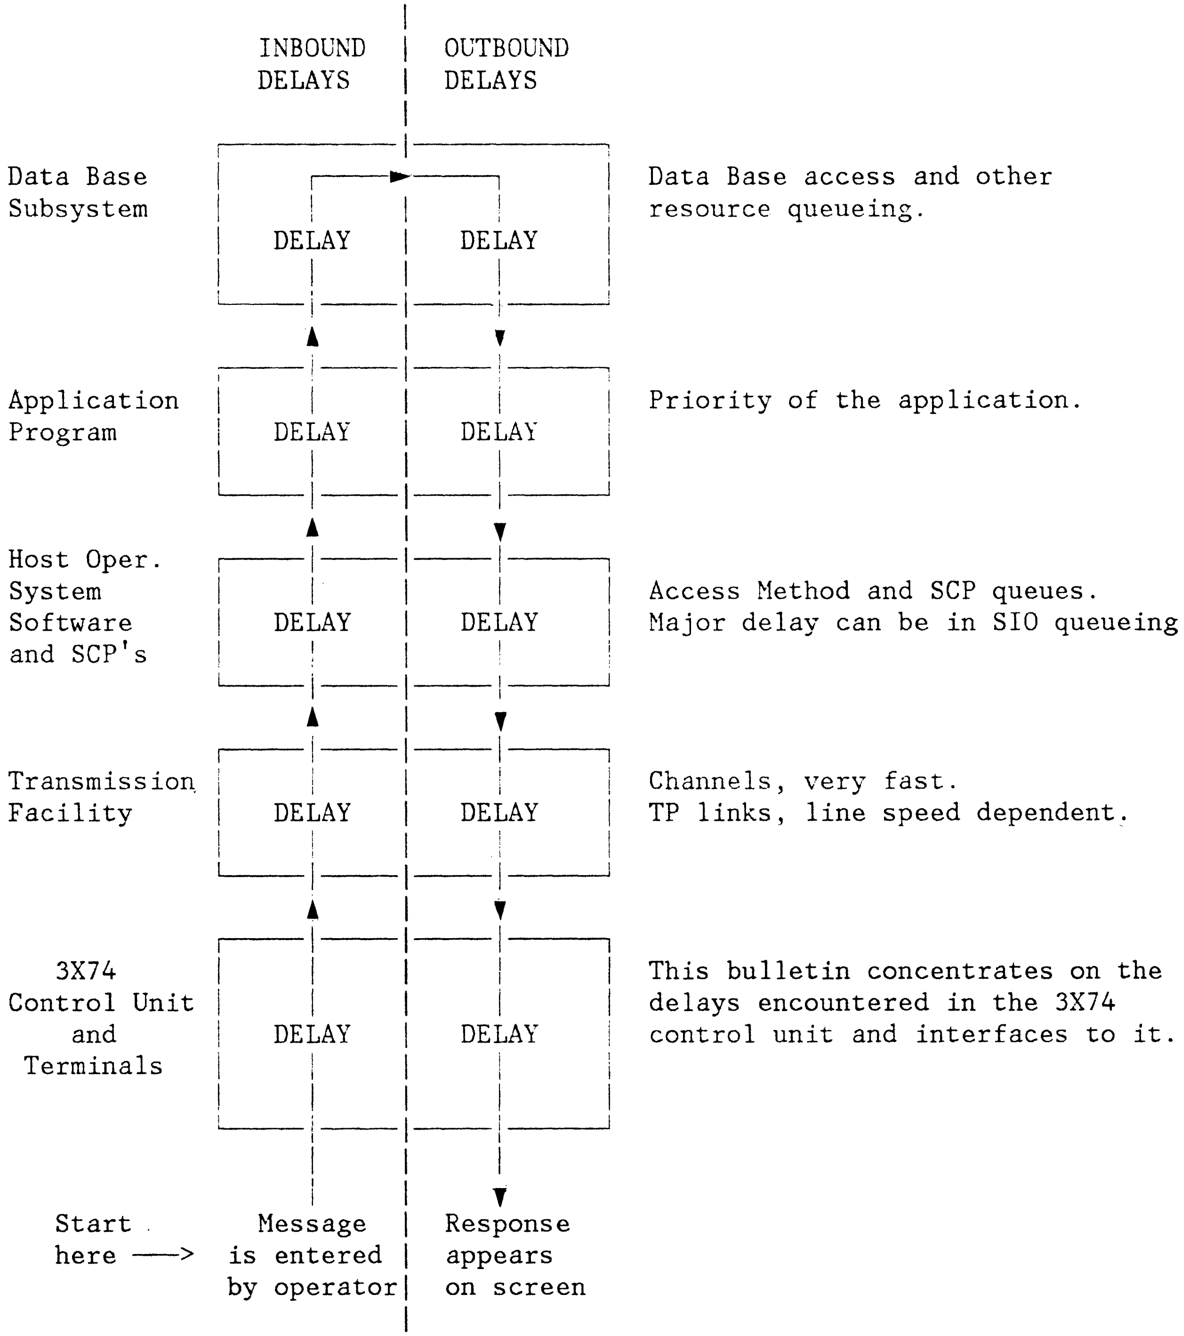

Figure 1. Delays in an Interactive System

Response times for transactions at a **remote** subsystem depend on the transmission rate over the telecommunication link. (See Chapter 9.) Outbound transmission time mostly overlaps 3X74 processing time, even for 64 kbps lines. For inbound messages, transmission time is added to the control unit time.

"Response", as used by others, sometimes implies response to **first** character. This is meaningful, for example, for judging the response capability of telecommunication networks because response to first character does not include the time needed to paint the screen, which may take relatively long, and is dependent on link speed and message length.

SUBSYSTEM CONTROL UNIT UTILIZATION

The utilization of a subsystem control unit is defined as the average message service time in the control unit multiplied by the average number of messages per second serviced by the control unit (multiplied by 100 to express utilization as a percentage). (In SNA control units, this is a slight oversimplification, because the moving and processing of RUs associated with different data streams can overlap.) High utilization of a channel-attached subsystem can increase its response times appreciably because of queuing delays being added to the (single-thread) service times of messages.

This average control unit utilization is a measure for the aggregate load that a local subsystem can handle with acceptable response times. As a general rule, it should not exceed fifty percent, or be less than twenty-five percent when maintaining minimal subsystem response time is very important. Note that for workstations operating in CUT mode, keystroke processing adds to 3X74 utilization; in DFT mode operation it does not.

Response times can also be affected by high **channel utilizations**. Normally, the contribution of data transfers over the channel to response time is small. However, when channel utilization is high, significant delays may be introduced. For this reason, one should plan to keep the average utilization of a channel by all its attached subsystems and devices below thirty percent. (See Appendix D.)

For remote subsystems, the speed and type of the communication link rather than the control unit determine performance, because control unit processing is faster and is overlapped by the data transfer over the link. For remote applications, it is therefore important to minimize data stream **length** to keep utilization of the link down.

In addition to selecting appropriate protocol parameters, data stream size may be reduced by invoking compression or compaction functions, as is available on some high speed printers, and by using appropriate 3270 data stream orders. For example, a row of eighty dashes across a page can be compressed to four characters by using the Repeat-to-Address order.

The capacity planning approach for 3X74 subsystems is based on adding the contributions of the various MFI transactions, concurrent file transfers, and printer operations to obtain subsystem utilization.

"OPEN" AND "CLOSED" LOOP OPERATIONS

The terms "open-loop" and "closed-loop" pertain to a characteristic of the transaction load.

When a transaction load is characterized as **open-loop**, it is assumed that this load is not affected by the response times of the host and the subsystem. For example, if a user at an interactive workstation generates an average of five transactions per minute, it is implied that this rate will be sustained regardless of the user-perceived response times being experienced, that is, it is "open-loop".

Another example of an open-loop load is the transaction load generated by a printer attached to a 3X74 control unit. Its demand for data is uniquely dependent on the printer/document characteristics, and should be independent of host and/or control unit 3270 data stream orders. (Depending on the size of printer data rate demand and host and/or control unit characteristics and loading, this demand may not always be met.)

For a **closed-loop** transaction load, there is a direct dependency of the transaction rate on the response time of all elements in the loop, that is, the host, the subsystem, and the operator/workstation. The operator/workstation and host elements may be assumed to have specific constant reaction times, so that the total time to complete a transaction will increase or decrease as the subsystem response time increases or decreases.

A file transfer operation is considered closed-loop because the host and a specific workstation are assumed to respond after fixed times. (For workstations, this fixed response usually differs for different types.) Therefore, if the load on a subsystem (and as a consequence its response time) increases, file transfer rate will slow down.

This differentiation between open-loop and closed-loop operations is a relative one, of course. Eventually, when a host or system response becomes large enough, any open-loop operation will become closed-loop.

IBM 3X74 SUBSYSTEM CONTROL UNIT MODELS

There are nine models of the IBM 3174 Subsystem Control Unit: 1L, 1R, 2R, 3R, 51R, 52R, 53R, 81R, and 82R. They are listed together with the older IBM 3274 models in Figure 2 on page 8, along with some of their principal characteristics. Certain functional and performance characteristics are obtained when combining one of these models with a given Configuration Support (microcode) level.

The location of the people usually mandates the location of their workstations, which, in turn, determines the manner in which subsystems are attached to a host. Since channel attachment is limited to 122 meters (400 feet) for the 3174, control units beyond that distance must be serviced by attachment to a telecommunications line (unless the IBM 3044, or other channel-extending hardware has been installed).

As many as four IBM 3299 Terminal Multiplexers Model 1 may be attached to a 3X74 control unit up to 1500 meters (4920 feet) away with coaxial cabling. In turn, eight terminals attach to a 3299 with up to 1500 meters (4920 feet) of coaxial cabling, allowing a maximum of 3000 meters (9840 feet) between terminal and control unit. Note that one 3299 permits as many as eight terminals to share as much as 1500 meters (4920 feet) of coaxial cable.

Using IBM Cabling System (ICS) type 1 or 2 wiring (or type 3 wiring in combination with up to four 3299 Models 2 or 3) may reduce the maximum cable distance of 1500 meters to 274 meters (900 feet), or somewhere in between, depending on the wiring configuration used.

Performance is often important in the selection of the attachment mechanism as well. The much higher data transfer rate available with host channels (in the order of one Mb per second or less) gives local control units a clear performance advantage over remote subsystems, where many consider 9600 bps a high speed link. For higher link speeds, up to 56 kbps for SNA, the gap narrows. This difference in remote versus local performance may dictate local channel attachment.

In local control units, it is the data stream processing rate in the subsystem rather than the channel transfer rate that determines subsystem response time. Because channel utilization by a single control unit is usually small, several control units can share a channel.

For local control units operating in an SNA environment, performance of the 3174 is considerably better than of the 3274 in terms of channel utilization, pass-through rate, and response time, with some exceptions as noted in the text. For remote subsystems, however, SNA/SDLC generally provides superior link performance compared to (non-SNA) BSC.

After having chosen a primary attachment method, some further determinations may need to be made, such as whether to operate a workstation in CUT or DFT mode (when the option exists), whether to trade off the number of attached terminals against performance, and the control unit type and model. The utilization of local subsystem control units is often a principal determinant of subsequent design activity, especially when file transfer, graphics, image, or a high speed printer is involved. When 3X74 utilization is kept within bounds, response time will not be unduly stretched out.

As in most IBM 3274 Control Units, data stream handling functions in the 3174 are implemented with a processor and microcoded logic to provide flexibility, easy functional growth and upgrading, etc. The microprocessor in the 3174 is about twice as fast as the one in 3274 Model 41s. Subsystem response does not decrease proportionally, however, because other subsystem elements affecting response time are identical, for example, the information rate in the cable between the workstation and the control unit.

The 3174 Models 1R, 2R, 51R, and 52R have about the same performance characteristics as the 3274 model 61C.

The 3174s are equipped with a faster microprocessor and other innovations which have improved their performance with respect to 3274 control units. For single messages, the performance differences between local control units operating in SNA and non-SNA mode are minor in comparison with other factors being addressed. Usually, only SNA performance data has been provided unless the subsystem will operate exclusively in non-SNA environments. For long messages and file transfers, significant performance differences may exist.

All performance data is based on using IBM 3299 Terminal Multiplexers, or the functional equivalent Terminal Multiplexer Adapter feature (#3101), and a total cable length of 150 meters (492 feet). Although there are additional propagation delays for longer distances, they can be ignored for the purposes of these guidelines. Because only one terminal is receiving or sending data at any one time, sharing of the single cable attaching a terminal multiplexer by eight ports will not introduce additional queues.

3174 Subsystem Control Units can attach to a host system through a host channel, telecommunications link, or IBM Token-Ring Network. The 3274 control units attach through a channel (local), or telecommunication link (remote) only.

Performance considerations for channel-attached control units are addressed in Chapters 3 and 4, including 3174/3274 performance-related differences. (See also **IBM 3174 Subsystem Control Unit: Host Programming Considerations,** GA23-0325.) The performance of 3174 subsystems incorporating an IBM Token-Ring Network is discussed in Chapter 10, and of telecommunication link attachments in Chapter 9.

Control Unit Model (they attach up to 32 devices, except as noted)	Protocol, Environment	Attachment	Storage[6], Kb 64	128	192	1024	1536	Configuration Support A	B	C	D	A[4]	S[5]
3174-1L	non-SNA/SNA	Channel				X	X[6]					X	X
3174-1R,-51R[2],-81R[3]	non-SNA/SNA	TP				X	X[6]					X	X
3174-2R,-52R[2],-82R[3]	non-SNA/SNA	TP				X	X[6]					X	X
3174-3R,-53R[2]	SNA	Token-Ring				X							X
3274-1A	SNA	Channel	M					X	X				
				X						X			
-21A	SNA	Channel	X					X	X				
-31A	SNA	Channel		X						X			
					X						X		
3274-1B[1]	non-SNA	Channel	X					X					
-21B[1]	non-SNA	Channel	X					X					
3274-1C/-51C[2]	non-SNA/SNA	TP	M					X	X				
				X						X			
-21C	non-SNA/SNA	TP	X					X	X				
-31C/-51C[2]	non-SNA/SNA	TP		X						X			
					X						X		
3274-1D	non-SNA	Channel	M					X	X				
				X						X			
-21D	non-SNA	Channel	X					X	X				
-31D	non-SNA	Channel		X						X			
					X						X		
3274-41A	SNA	Channel			X						X		
3274-41C/-61C[2]	non-SNA/SNA	TP			X						X		
3274-41D	non-SNA	Channel			X						X		

NOTES:
- X Normal use
- M Minimum memory size, 128 Kb may be required
- [1] Data stream processing performed by hardware
- [2] These models attach up to 16 devices
- [3] These models attach up to 8 devices
- [4] Configuration support for 3174 control unit
- [5] Configuration support for 3174s with Token-Ring Network
- [6] Up to 3 Mb, in 0.5 Mb increments

Figure 2. IBM 3174/3274 Model Characteristics

CHAPTER 2. CONTROL UNIT FUNCTION IN CUT AND DFT MODES

Workstations attached to a 3X74 control unit share it for performing some or all of the following functions:

- Linking the subsystem to a host system through a channel or telecommunication line
- Processing keystrokes entered at the workstations
- Executing outbound data stream commands and orders to expand and otherwise prepare data for updating of a workstation's display screen
- Preparing inbound data streams for transfer to the host in response to host commands and workstation buffer contents, for example, keyed data
- Linking displays and printers for COPY operations.

The allocation of functions, or 'function split', between subsystem control unit and terminal depends on the workstation type and mode of operation. In Control Unit Terminal (CUT) mode, the control unit performs all of the above, while in Distributed Function Terminal (DFT) mode, the workstation does all keystroke and data stream processing. These modes exhibit different performance characteristics.

- **Control Unit Terminal (CUT) mode:** The control unit performs all listed functions, including processing of keystrokes and data streams. Data stream service times depend on their length and content, the speed of the processor, the microcode, and the delays incurred in accessing data items in the buffer of the attached device. Some terminals, such as the IBM 3278, will display the last character on the screen as soon as the control unit has processed it. In others, such as the IBM 3180, there may be a small delay, depending on workstation design.
- **Distributed Function Terminal (DFT) mode:** The control unit passes data streams back and forth between host and workstation. The workstation executes or prepares data streams, and also processes keystrokes. Examples are the IBM 3290 Information Panel, the IBM 3192 G Color Graphics Display Station, and the 3193 Image Display Station. For a given control unit and data stream, service time tends to be shorter than for CUT mode, and is dependent on data stream length only. The workstation service time depends on data stream length and content, the speed of its processor, the effectiveness of the microcode, and the amount of display screen processing.

Although pass-through time for the control unit in DFT mode may be less than for CUT mode processing time, this gain is often reversed by the processing time in the workstation, resulting in longer subsystem response time in DFT mode.

A positive effect of DFT mode on response time is the off-loading of the control unit by the attached device. Therefore, control unit utilization and average queueing time build more slowly with increasing transaction load than in CUT mode, allowing the subsystem to handle larger transaction loads (more workstations, file transfers, and printers), and/or reduce queueing delays. (See Figure 3 on page 10.)

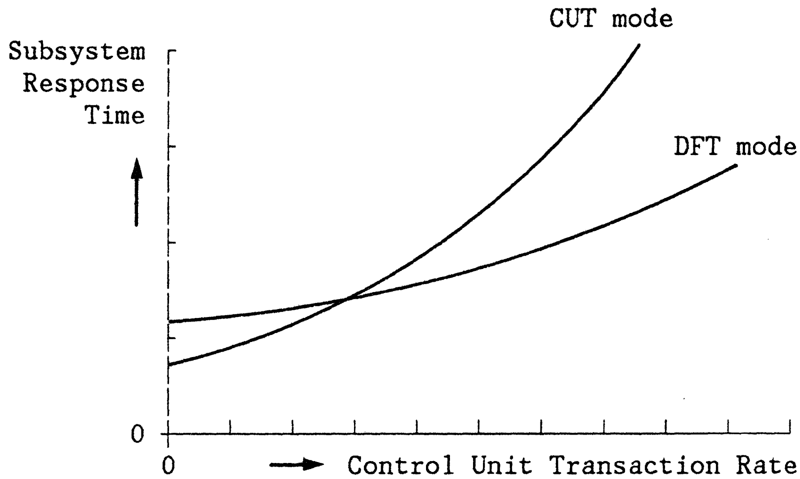

Figure 3. CUT and DFT Mode Comparison

The IBM 3178, 3179, 3180, 3191, 3192, 3278, and 3279 Display Stations exemplify CUT mode operation, while the IBM 3290 Information Panel and the IBM 3193 Display Station exemplify DFT mode operation.

Summarizing for current equipment, CUT mode will often yield the fastest response times for lightly to moderately loaded subsystems, while in DFT mode control unit load, and therefore response time, will tend to increase less with increasing transaction load. For this reason, when there is a choice between CUT and DFT mode, the selection of DFT mode is recommended unless there it is very important to obtain the shortest possible subsystem response time.

With many DFT workstations it is possible to emulate up to four logical terminals, and conduct a host session with each. This requires the assignment of several terminal addresses to the control unit port to which the workstation is attached (by customization of the control unit microcode).

The use of DFT mode in 3274 Control Units requires configuration support D, Release 64.2 or higher; earlier 3274 support and remote 3276 control units support CUT mode operation only. On both the 3174 and 3274, the customization procedure can only be executed on a CUT display station connected to port 0. When not used for this purpose, port 0 can support a single session in DFT mode.

MULTIPLE LOGICAL TERMINALS OPERATION

In Multiple Logical Terminals (MLT) operation, an operator uses a single workstation with one keyboard for conducting up to five host sessions simultaneously. This capability has existed in several DFT display stations for some time, for example, in the IBM 3290, but is now also available to users of CUT terminals.

Single-session applications with emphasis on productivity have been assumed to operate at an average rate of 6.7 transactions per minute, or less, depending on the application. The principal purpose of this section is to consider modification of this assumption where necessary.

Because operational and MLT performance characteristics for CUT and DFT mode differ somewhat, they are addressed separately.

DFT MODE

When a workstation, for example, a 3290, is logged on to more than one session, the screen for each session is displayed.

In other DFT workstations, such as a Personal Computer AT, only the screen for the active session may be displayed. When the operator moves to (activates) another session, the workstation displays the buffer associated with the newly activated session. The display buffers for all sessions are always present in the workstation, and kept up to date.

The transaction load in the control unit may increase when the operator initiates transactions in rapid succession by frequently transferring the cursor from one session to another. In theory, this could result in a queue of up to four outbound messages in the host.

However, DFT workstations configured as two, three, or four terminals are not likely to generate two, three, or four times the transaction load because there is only one keyboard and one operator. Depending on the application and number of concurrent sessions, use a somewhat increased transaction rate, but not to exceed fifty percent, for such a workstation.

The same rationale can be used for host programs creating multiple partitions on a display screen. As with the MLT function, the use of slightly increased transaction rates should sometimes be considered, because interleaving of operator transactions may occur.

CUT MODE

When a CUT mode workstation is used for MLT operation, the display buffers associated with the sessions reside and are maintained in the 3174 control unit. Up to five concurrent sessions can be established, which can be accessed in round robin fashion by actuating a change-screen key sequence.

In the MLT environment, the 3174 service time and utilization of a type A-1200 transaction, and of most others, do not exceed the single-session times by more than two percent, on the average.

Every time the operator uses the change-screen key sequence, up to about 70 milliseconds of 3174 time may be used which includes changing the 1920-character screen. This is less than the control unit utilization by a type A-1200 transaction.

To estimate the 3174 utilization increase as a result of an operator conducting several (up to five) MFI sessions simultaneously, consider the following. To randomly select another one of five sessions, an operator will have to actuate the change-screen key sequence an **average** of 2.5 times because of the round robin selection approach, that is, the average of one, two, three, and four times. (This average with four, three, and two sessions is 2x, 1.5x, and 1x, respectively.)

Counting a change-screen key operation as a transaction, you can estimate control unit utilization by adding your estimate for the number of times per minute the operator will use the change-screen key to your estimate for the number of transactions that will be initiated per minute.

Because there may not be as much 'think time' associated with actuating a change-screen key, one would expect the transaction rate on an MLT workstation to increase somewhat over a single-session transaction rate, but not be a multiple.

PROGRAMMABLE WORKSTATION OPERATIONS

For some programmable workstations, such as the IBM 3270PC/AT, the user determines whether to interact with the control unit in CUT or DFT mode (provided the control unit has been customized for DFT mode).

Whether in CUT or DFT mode, these workstations may operate in one of three ways, sometimes concurrently or in combination,

- **As a stand-alone processor,** not communicating with the IBM 3270 control unit except for responding to polls, and, possibly receiving unsolicited outbound messages
- **As a host-interactive terminal,** emulating a 3270 display device
- **As a recipient or sender of files** of various lengths to or from a host (using a file device in the workstation).

The effect of programmable workstations on control unit utilization differs markedly, ranging from stand-alone mode with virtually no demands, to file transfer operations with the potential of high control unit utilization.

CHAPTER 3. IBM 3X74 CONTROL UNITS

In local IBM 3X74 subsystems, performance is largely determined by the control unit with its attached devices, rather than the channel transfer rate (up to 1250 kbytes per second for the 3174, less for the 3274).

Although the effect of channel rate on subsystem performance is minor, subsystem operations do add to channel utilization, and may therefore affect the total load that can be processed by a host channel. (See Appendix D.)

Early 3274 control units attached display stations and printers in CUT mode only (except for printers using the LU type 1 protocol). MFI data streams were designed to write and read screens of 960, 1920, 2560, 3440, and 3564 characters. Later, these sizes, except 960, were used in the IBM 3290, the workstation that introduced DFT mode.

With graphics, image, and file transfers applications, long data streams are now being used with other devices than just LU type 1 for printers.

For CUT mode, in both SNA and non-SNA environments, long data streams are divided in screen-sized amounts which are transferred in separate, mutually independent operations, as in non-SNA and LU type 3 printer data streams. Most programs transferring files between a host and an intelligent workstation divide a file into 2 or 3.5 Kb sections which are transferred and processed as individual 'screens' of data.

In DFT mode, data is passed through the control unit, without being examined, for processing in the workstation. For long data streams, this offers the opportunity for overlapping transfer through the control unit with processing in the workstation. Considerations associated with long data streams are being addressed in Chapter 4 on pass-through of long data streams in DFT mode.

A significant difference between CUT and DFT mode with respect to control unit utilization is that in CUT mode a control unit processes keystrokes, while in DFT mode the workstations relieve it from this demand. See the section on keystroke rates in CUT mode workstations on page 18.

RU/BLOCK SIZE SELECTION

The selection of RU (Request/Response Unit, in SNA) or block size (in non-SNA) affects subsystem performance and control unit utilization. This selection depends on the characteristics of the control unit and the attached device.

In SNA, data is transferred in RUs consisting of the data and an RH (Request/Response Header). The maximum RU length (not including headers) is specified by the host for each 3X74-attached device in the SNA bind when the LU-LU session is established. If the data stream length is less than this RU size, a single RU of data stream length (plus some bytes for headers) is sent. If the data stream is longer, more than one RU is used.

A 3174 permits up to 2048 and 4096 bytes total per RU for inbound and outbound data streams, respectively; up to 1024 and 1536 bytes for 3274 control units.

The RUs in a data stream may be paced, that is, a host may send one or more RUs as specified by the pacing count, when the receiving device returns a pacing response. The number of RUs that can be sent per pacing response is called the window. See Chapter 4 for the use of pacing to effect overlap of transfer with processing in the workstation.

In non-SNA, data is transferred in blocks with a maximum size of 15,600 bytes in 3174 control units, and 7168 bytes for 3274s. The maximum block size is specified in the host, and is used to divide a data stream in blocks if its length exceeds this maximum block size. The transfer of each block is an independent operation.

There is no equivalent of pacing in non-SNA, even for pass-through of long data streams; the host sends the next block in response to a "device end" (DE) returned by the subsystem, or an acknowledgment message from the receiving device.

In both SNA and non-SNA, CUT and DFT mode, maximizing the amount of data per RU or block enhances performance because it minimizes the amount of processing of headers, responses, etc. (overhead) for a given quantity of data. Data streams in MFI applications tend to be relatively short inbound (usually somewhere from 10-100 bytes), and up to about 2 or 4 Kb outbound. It is recommended to fit them, whenever possible, in a single RU or block.

For best performance of file transfer operations, for which files are usually divided in 2 or 3.5 Kb sections (in CUT mode, a workstation screen size), use of a single, or minimum number of RUs or blocks, is recommended as well.

3274 MODELS 21A, 31A AND 41A SNA CONTROL UNITS

As noted before, the 3274 Models 21A, 31A and 41A SNA control units support maximum inbound and outbound RU-sizes of 1024 and 1536 bytes, respectively. Their line buffer capacity is 4500 bytes. The line buffer is actually a pool of buffers, each of which is set to the lesser of 256 data bytes, or the value specified by bytes 6,7 (size of host buffers) of the Connect function of the channel Control command.

This Connect parameter corresponds to the amount of data that may be sent inbound with each channel Read command. In ACF/VTAM this is the value of BUFSIZE for the I/O buffer pool. It is a universal system parameter, and cannot be assigned different values for individual system components. Since the 3274 inbound and outbound processing is by line buffer, performance will be improved if this parameter is set to at least 256.

Data is transferred inbound by the channel Read program. This consists of a sequence of Read CCWs (commands) chained together. Each Read transfers the data specified by the Connect function. The number of Read CCWs in a Read program is specified by bytes 4,5 in the Connect function of the channel Control command. In ACF/VTAM this is the value of MAXBFRU. Since there is 3274 overhead

for each Read program, the performance will be improved if the value of this parameter is more than one.

If the line buffer size is set to hold 256 data bytes, then each buffer is assigned 272 bytes to accommodate the data, plus the TH and RH and internal 3274 headers. There are 16 buffers. SNA architecture supports interleaved inbound and outbound operations so these buffers are assigned dynamically to inbound or outbound, but a minimum of one fourth are reserved for inbound and one half for outbound, with the rest assigned FIFO. Therefore, there is at least 1 Kb of data available for inbound, and 2 Kb of data for outbound. The maximums are 2 Kb (data) for inbound, and 3 Kb (data) for outbound.

When a pacing count of two or more is specified, the length of all RUs sent as a result of a pacing response should not exceed 3 Kb outbound, or the subchannel will be kept busy until the 3274 can pass data through to the terminal.

Inbound length is not a concern since the 3274 effectively slices the total inbound data into RU-sized (max = 1024) pieces, and initiates channel transfer by asynchronous interrupt as each RU becomes available in the inbound buffers.

3174 MODEL 1L SNA CONTROL UNIT

The 3174 Model 1L control unit, when customized for SNA, supports maximum inbound and outbound RU-sizes of 2 Kb and 4 Kb, respectively.

This SNA control unit has an aggregate of 17,420 bytes of (line) buffers for storing inbound and outbound messages being sent, or received, over the channel. They are permanently partitioned as follows:

Outbound: 5 of 2K, and 27 of 100 bytes each
Inbound: 4 of 1K, and 8 of 48 bytes each

Short messages will automatically use the small buffers. The larger inbound and outbound sizes are used for longer messages. The additional SNA header bytes are also accommodated. The Connect function does not effect the size of the 3174 buffers which are fixed. An RU can span two buffers but only one RU will be assigned to a single buffer.

Since the data is processed by buffer, performance will improve when larger RUs are used. For outbound operations, the number of large buffers required to accommodate the RU, or RUs, sent per pacing response should not exceed five, to avoid excessive channel utilization. Inbound is not a concern since the operation is by asynchronous interrupt as in the 3274.

The same channel Read Program recommendations for the 3274 apply to the 3174.

3X74 NON-SNA CONTROL UNITS

The 3174 Model 1L, when customized for non-SNA, has line buffers totaling 15,612 bytes; 7168 bytes in the 3274 Models 31D and 41D. To minimize channel utilization, the amount of data per block should not exceed this limit.

In these non-SNA control units, the entire line buffer is dedicated to a single inbound or outbound transaction.

These line buffer sizes are adequate to accommodate the data streams associated with most alphameric MFI transactions. Long data streams, as may be found in graphics and image applications, need to be sectioned in blocks not exceeding this limit. (See Chapter 4.)

KEYSTROKE RATES IN CUT MODE WORKSTATIONS

For workstations operating in CUT mode, keystrokes are processed in the 3X74 control unit. Keystrokes are queued in the 3X74 to await service, one queue for each terminal. A keystroke is moved into its queue within a few milliseconds after the key has been pressed.

The rates at which IBM terminals may generate keystrokes are:

- In typematic mode, at 10 cycles per second. (The horizontal cursor-move keys on some recently announced terminals generate strokes at 20 cps.)
- An average of 5 cps in manual typing mode by a skilled typist.
- 10 cps in playback mode of previously recorded keystrokes.

This keystroke traffic is concurrent with host interaction and other activity in a 3X74. Because servicing of keystrokes has a higher priority than most other functions, the 3X74 capacity limit for MFI and file transfer operations is usually reached before keystroke performance degrades. In general, keystroke performance has not been a concern when assessing 3X74 performance.

If the rate at which keystrokes are generated exceeds the rate at which the 3X74 can process them, then the keystroke queues will start to fill and eventually overflow. Failure of the control unit to keep up will result in either or both of the following conditions:

1. A perceptible lag between key actuation and the appearance of a data character on the screen, or movement of the cursor.
2. Keyboard lock, and the appearance of X?+ in the operator information area on the screen. (Pressing the Reset key will cancel this indicator.)

The first condition will occur as a keystroke queue starts to fill. The second condition results from a queue overflow; it can be forced by pressing many keys simultaneously. (Other events, not addressed here, may also invoke the X?+ indicator.)

Other 3X74 activity, for example, interaction with the host, may use the microprocessor and/or the coax adapter, thereby reducing their availability for servicing keystrokes. Past experience with and/or analysis of 3270 applications has shown that the 3X74 control units can keep up with the keystroke traffic generated by IBM 3270 CUT mode terminals.

A limited experiment with a number of different applications in a professional environment yielded an average 3174 utilization by keystroke processing of about 0.3 percent for an active CUT workstation used at a rate of 6.7 transactions per hour. (This compares with about one percent for data stream processing and transmission at the same transaction rate.) This may well be higher for some data entry in a production environment, and lower for inquiry applications.

When using workstations with typematic or playback rates in excess of 10 cps, there is an increased probability that one or both of the two failure modes will occur. The 3274 control units (especially Models 31 and earlier) are more susceptible than 3174s because they have less powerful microprocessors. For example, a 3274 Model 31C on a 4800 or 9600 bps line is most likely to give problems because the processing of BSC link data represents a considerable load for the microprocessor with a higher priority than keystroke processing.

The probability of failure is difficult to predict because it depends on many variables, such as the number of active CUT mode terminals, duration and type of keystroke activity (application), link speed, and link utilization by MFI transactions and others (file transfer). Experience has taught that a 10 cps keystroke limit on typematic and playback provides safe operation. With a higher limit it may be necessary to reduce the number of attached terminals to achieve good keystroke performance.

DFT WORKSTATION RESPONSE TIMES WITH SHORT DATA STREAMS

For DFT mode with short data stream applications, such as alphameric MFI, the 3174-1L contribution to subsystem response time is reduced by 10-25 percent relative to that of the 3274-41 control unit. This improvement may not always be noticed by the user because many DFT workstations process a data stream more slowly than the control unit passes it through, and are therefore responsible for the major portion of the subsystem response time. However, control unit utilization per transaction is reduced, and the control unit can support a larger transaction rate.

The service times in Figure 4 on page 18 are for 10 bytes of inbound or outbound data, and do not include DFT device or host CPU delays. They represent relative figures of merit for the control units.

An inbound operation may be the pressing of the Enter key after keying in a CMS command, or using a PF key for paging through a file. The outbound operation might be a system acknowledgment to a command or RSCS status.
Non-SNA overhead is higher than SNA because the 3X74 must decode device dependent channel commands and translate them into terminal communications adapter (TCA) interface function requests which it puts in the TCA buffer control space for execution by the terminal. For SNA, the terminal commands are contained in the RU which the 3X74 passes through to the terminal Terminal Communications Adapter (TCA) buffer without inspecting.

Control Unit	SNA Inbound	SNA Outbound	Non-SNA Inbound	Non-SNA Outbound
3174	22	20	31	30
3274-41	15	20	28	31
3274-31	25	35	44	47
NOTE: Pass-through times for 10 bytes, in milliseconds				

Figure 4. DFT Pass-Through Times for Short Data Streams

SNA PERFORMANCE OF 3174 VERSUS 3274

Comparing 3174 performance with that of the 3274 in SNA environments:

- The 3174 maximum RU sizes are larger than of the 3274, making it possible to reduce overhead processing for larger data streams. (Outbound 4 Kb versus 1.5 Kb; inbound 2 Kb versus 1.0 Kb, respectively.)

- Lower channel utilization as a result of higher 3174 channel transfer rates. (About one Mb per second versus 0.1 Mb per second for the 3274.)

- Because the 3174 has a total of 16 Kb of channel buffers versus 4.5 Kb in the 3274, the 3174 channel buffers are less likely to be depleted in high traffic situations than those in the 3274, resulting in reduced performance restrictions (Included in the 3174 buffers are twenty-seven 100-byte outbound buffers and and eight 50-byte inbound buffers. These are used for short messages. The remainder of the 3174 buffering is for long messages; 10K outbound and 4K inbound.)

- The 3174 DFT pass-through rate is two to three times that of the 3274. The overall throughput capacity of the 3174 is about two times that of the 3274.

- Alphameric CUT mode response times are about 50 percent of the 3274's.

- The internal 3174 structure (microprocessor and adapters) is more efficient. Charts 6-1A, -2A, and -3A (CUT mode) clearly demonstrate that, as traffic rates increase, the 3174 advantage over the 3274 increases.

SNA VERSUS NON-SNA PERFORMANCE

Comparing the use of the 3174 in SNA and non-SNA environments:

- The 3174 performance in SNA environments equals, or is slightly better than its performance in non-SNA environments. (For the 3274, SNA performance was worse than non-SNA performance due to the slower channel interface.)

- Comparing charts 6-1A and 6-1B show that in multi-thread, high traffic situations, the 3174 SNA advantage over non-SNA increases. The reason is that non-SNA architecture is single-thread, while SNA allows RUs to and from different workstations to be processed simultaneously.

CHAPTER 4. PASS-THROUGH OF LONG DATA STREAMS IN DFT MODE

Long data streams, as may occur in graphics, image, and LU type 1 printer applications, must be divided into many RUs (in SNA environments) or blocks (for non-SNA) when transmitted, in order to be accommodated in the buffer resources in the control unit and, for DFT mode, in the workstation as well. This chapter discusses what can be done to obtain their optimum performance. (In CUT mode and file transfers these issues hardly arise because the subsystem handles the relatively short data streams associated with their operations individually.)

The key to maximizing pass-through rates of long data streams in DFT mode is overlapping the operations in the control unit and the workstation. In SNA, this is possible by taking advantage of pacing, which permits more than one RU of a chain to be in transit between the host and a DFT workstation. In this way, an RU can be available in the terminal communications adapter (TCA) buffer for processing, as soon as the workstation has finished processing the previous one. (The same principle applies to LU type 1 printing.)

In non-SNA, there are also options that permit operations to overlap, for example, early print control (EPC) for printers. (See Chapter 11.)

BUFFER SIZE IN DFT WORKSTATION ADAPTER

The size of the TCA buffer in the DFT workstation (or of the buffer in a printer adapter) constrains the size of the RU or blocks in which long data streams are divided. If this buffer is too small to accommodate the data, buffer resources in the control unit will be used, with possible adverse effects on the performance of the subsystem, or even of the channel.

The size of the TCA buffer in most workstations is 8 Kb, including those using the 3278/79 Emulation Adapter (#5050). (Early versions of the 3290, and the #2507 3278/79 Emulation Adapter, now superseded by #5050, have a 4 Kb buffer in the TCA.)

About 7.5 Kb (3.5 Kb) is available for data stream storage, because 0.5 Kb is used for control blocks.

DFT AND LU TYPE 1 PASS-THROUGH OF LONG DATA STREAMS IN SNA

For the efficient transfer of long data streams through 3X74 SNA control units in DFT mode, the selection of RU size, pacing, and response ("exception" or "definite") matters because it takes many RUs to transmit them.

For 3274 Models 31A and 41A control units, make sure that **RPQ #8K1311, DFT Display Perform Enhance,** is installed for achieving optimum SNA performance. (This function is standard in the 3174 microcode.) Without this RPQ, outbound throughput rates may decrease as much as thirty percent, depending on the VTAM buffer (segment) size. (256 bytes, for example, would increase response time

less than 128 bytes.) For alphameric applications with short data streams (usually < 2 Kb), the benefit of the RPQ will be small.

For long data streams, the use of pacing is necessary to avoid overflow of the TCA buffer in the workstation and the line buffers in the control unit.

The pacing count (=N) controls the number of RUs that can be sent in response to a pacing response from the receiver. The RUs must belong to the same chain, which is the reason why pacing has little meaning for data streams with few RUs.

The amount of outbound data sent per pacing response plus data remaining from the previous transmission should not exceed the size of the DFT device buffer. Furthermore, the amount of data sent per pacing response is constrained by the size of the outbound 3X74 line buffers. The data length in the TCA buffer is determined by:

(2N - 1) x RU size bytes

For example, with N=2 and a TCA data buffer size of 7.5 Kb, an RU size equal or less than 2.5 Kb should be chosen. This would be acceptable for a 3174 control unit, but not for a 3274 model 41A with a maximum RU size of 1536 bytes. With an RU size of 1536 bytes and a 7.5 Kb data buffer, a maximum pacing of N=3 could be selected for the 3174. (N=2 for the 3274 because its maximum line buffer size is 3 Kb.)

The formula implies that as soon as the workstation has processed the first RU (of N) in the TCA buffer, there is enough buffer space to receive N more RUs from the host ((N-1)+N=2N-1), and that therefore a pacing response can be returned to the host. Thus, data transfer through the control unit and processing in the workstation overlap!

The words "the workstation has processed" need some elaboration. Some workstations process an RU while resident in the TCA buffer, so that the pacing response can only be sent after processing is complete. In other workstations, the RU is first moved to another storage area and then processed. Because this frees up its TCA buffer space, the pacing response can be sent sooner, and the amount of overlap is increased. Recovery is not a problem because the RU can be identified to the sender.

For N=1 the formula allows the maximum RU size to be 7.5 Kb but is limited, of course, to 4096 bytes in the 3174 and to 1536 bytes in the 3274. In this case, pass-through and processing in the workstation do not overlap (in the case that RUs are processed from the TCA buffer). The workstation is idle while the host prepares the next RU and passes it through the control unit.

The use of "exception" rather than "definite" response is recommended, because it reduces traffic through the 3X74.

From Chapter 3 we know that the 3174 Model 1L has 10 Kb of outbound buffers in 2 Kb partitions. Any overflow of an RU from one 2 Kb partition will use a second 2 Kb partition. This suggests that it is probably better to use 2 Kb rather than 2.5 Kb for better utilization of the 3174 buffer resource. (The 3274 has a 3 Kb outbound buffer pool from which storage is allocated as required.)

Data Rates, in kilobytes per second									
SNA		Req'd TCA	Outbound			Inbound			Recommended for 3X74/TCA buffer size combination:
RU		buffer	3274		3174	3274		3174	
size	Pac	Kb	-31A[1]	-41A[1]	-1L	-31A	-41A	-1L	
256	1	0.5	6.2	9.9	12.0	6.6	8.7	10.6	
512	1	0.5	10.1	15.5	22.4	8.4	10.7	18.7	
1024	1	1	14.7	21.6	37.3	9.8	12.2	30.9	
1536	1	1.5	17.4	24.9	50.5	NA	NA	39.8	
1024	2	3	18.1	25.8	49.3	10.0	12.4	38.8	3X74/3.5
1536	2	4.5	20.4	28.4	65.3	NA	NA	48.2	3274/7.5
2048	1	2	NA	NA	58.5	NA	NA	46.9	
2048	2	6	NA	NA	73.4	NA	NA	55.5	3174/7.5
4096	1	4	NA	NA	74.5	NA	NA	NA	
4096	2	12	NA	NA	87.0	NA	NA	NA	

NOTES AND ASSUMPTIONS:

[1] With RPQ 8K1311, DFT Display Performance Enhance, installed.

- NA - Not applicable
- Delays in host and DFT workstation for data stream, Enter key and function request processing, TCA handshaking, pacing response, etc. not included, must be taken into account for total systems performance.
- A 500 foot coaxial cable length is assumed. (For 3274s, longer cables (up to 10,000 feet) and a 3299 may reduce these rates up to about 2 5 and 7 percent for inbound and outbound, respectively. For 3174s, rates are not significantly degraded; connection through a 3299 or TMA included.)
- Pass-through rates are based on time to pacing response.
- Inbound rates include Read; outbound rates include Write.
- Data transmission is not held up by processing or a shortage of buffer space in the workstation.

Figure 5. DFT Pass-Through Rates for SNA Control Units

The table in Figure 5 on page 23 lists pass-through rates, in kilobytes per second, on the basis of data sent per pacing response. For example, "RU=1536 with pacing =2" means that two RUs of 1536 bytes, that is, a total of 3072 bytes, are sent per pacing response. Data streams longer than 3072 bytes would repeat the sequence until completion. The amount of data that should be sent per pacing response is limited by the 3X74 buffers. A lower limit may be imposed by the TCA buffer in a DFT workstation. These rates in do not include DFT device or host times.

The best achievable control unit performance is given here. Reasons for these rates being much lower than the instantaneous transmission rate of 196 kbytes

per second in the cable to a workstation are overhead in the 3X74, moving data within the control unit, polling for keystrokes, etc.

For the 3274, inbound rates are small relative to outbound. A major reason is that DCA Read Multiple is not used by the 3274 microcode.

The improved 3174 pass-through rate with respect to the 3274 Model 41A rates are mostly due to the new channel adapter design in the 3174.

In summary, for best performance of applications with long data streams, use a pacing count of at least N=2, and the maximum RU size as allowed by either the control unit type, the control unit line buffer size (equal or larger than NxRU, and, for the 3174, with awareness of the 2 Kb granularity), or the TCA buffer length (equal or larger than 2N-1). With a maximum amount of data in a single RU, the number of RUs needed to transfer a long data stream, and thereby the overhead through all components of a system, is minimized. They can be set independently in the bind for each logical unit (LU) connected to a 3X74.

DFT PASS-THROUGH IN NON-SNA

In non-SNA DFT mode, single blocks of data are passed through the control unit, with each block being an independent transmission. This operation resembles the pass-through of RUs with pacing N=1.

To minimize processing overhead, it is recommended to maximize block size by taking advantage of the large block sizes allowed in non-SNA control units (15,612 bytes for 3174, and 7168 bytes for 3274). Exceeding these lengths will increase channel utilization. In non-SNA, the line buffers in the control unit are used by only one inbound or outbound transaction at the time during pass-through of the inbound and outbound data streams of an MFI transaction. The control unit is available during the interval that the DFT workstation processes a data stream.

The 7.5 Kb TCA data buffer in almost all current DFT workstations makes 7168 bytes a universal limit at this time, and is the maximum block size in use for image and graphics. GDDM, under which graphics and image run, supports a 7168 byte block. (The default is 1536 bytes.)

Note: Because there are some older workstations with a 4 Kb TCA buffer, and there is no equivalent of the 'bind' in non-SNA, one may have to use a smaller block size, even though some or most of the DFT workstations may have an 8 Kb TCA buffer.

The raw data rates for non-SNA as shown in Figure 6 on page 25 are, in some cases, higher than for SNA but at the system level this advantage may disappear. As for SNA, the best achievable throughput rate is listed. For similar reasons, these rates are much lower than the instantaneous rate of 196 kbytes per second in the cable connecting a workstation.

As noted, non-SNA operation can be likened to an SNA operation with pacing N=1. The only possibility for achieving some overlap between workstation processing and data transfer is to return DE before processing of a block is completed. This may affect recovery procedures, however, because an error in a block for

Data Rates, in kilobytes per second						
Non-SNA	Outbound			Inbound		
Block size	3274 -31D	3274 -41D	3174 -1L	3274 -31D	3274 -41D	3174 -1L
256	5.3	7.7	8.3	5.3	8.0	7.7
512	10.0	14.3	15.9	9.8	14.2	14.7
1024	18.1	25.1	29.5	16.8	22.9	25.6
1536	25.0	33.7	39.5	21.6	28.4	34.1
2048	30.6	40.4	49.4	25.7	32.7	40.9
3584	43.8	55.0	68.5	33.0	39.8	54.9
7168	56.1	67.7	88.6	39.5	45.4	71.1
15612	NA	NA	104.6	NA	NA	84.2

NOTES:

For inbound, assumed DFT device supports NDCA Read Multiple 32 with the 3174.

For non-SNA, inbound rates include Enter, Select RM, RM, and RM; outbound rates include Write Structured Field.

Figure 6. DFT Pass-Through Rates for Non-SNA Control Units

which DE was returned prematurely cannot be identified to the host at a later time.

For non-SNA, the subject of 3X74 utilization for DFT workstation operation is addressed in Appendix C.

EFFECT OF PASS-THROUGH ON OTHER WORKSTATIONS

Pass-through of long data streams can interfere with the processing of data streams for CUT devices. The effects of long DFT data streams on CUT mode performance are shown in Charts 4-1 through 4-4.

Each chart has a family of curves showing the effect of various DFT mode transfer rates on 3X74/3278-2 response times for the A-1200 benchmark. The curves labeled "0" show the response time versus A-1200 transaction rate without DFT pass-through. The other curves show the response time with increasing DFT transfer rates, in RUs per minute for SNA, and in blocks per minute for non-SNA. Curves are truncated prior to the right margin at the point where A-1200 transaction rates cause control unit utilization to exceed 65 percent.

The charts are used as follows:

1. Convert the DFT device processing rate into blocks or RUs per minute. For the performance of DFT graphic applications, see the **Graphic Systems Performance Guide,** ZZ20-5571, by A. J. Kirkland (IBM Hursley). It contains data lengths and transmission times for a number of graphics benchmarks. Note that some graphic displays have a separate draw time. This time should not be used since it occurs off line from the control unit. Use just the transmission time where there is a separate draw time, and transmission plus draw time where they are combined. For the 3193 image display, use the decompression rate provided in its data sheet.

2. There are curves for several RU and block sizes. Most applications will use one of these. Use the curve for the RU or block size that is closest to the RUs or blocks per minute value, or interpolate between curves.

3. Read the CUT mode response time corresponding to the total A-1200 transactions rate in the control unit.

This CUT mode response time is the mean value that can be expected during the time that data is being sent to a DFT device.

To get the time duration during which interference will occur, divide the data stream length by the data rate. For graphics and image data lengths, consult the sources referenced in item 1 above.

The frequency with which DFT interference occurs will depend on the number of active DFT workstations and their graphics or image transaction rates. Having more than one DFT device on a control unit may result in two or more simultaneous DFT transactions. This would increase the RU or block rate through the control unit proportionately, and possibly result in significant variability in CUT device response times. This should be taken into account when determining the mix of DFT and CUT devices on a control unit.

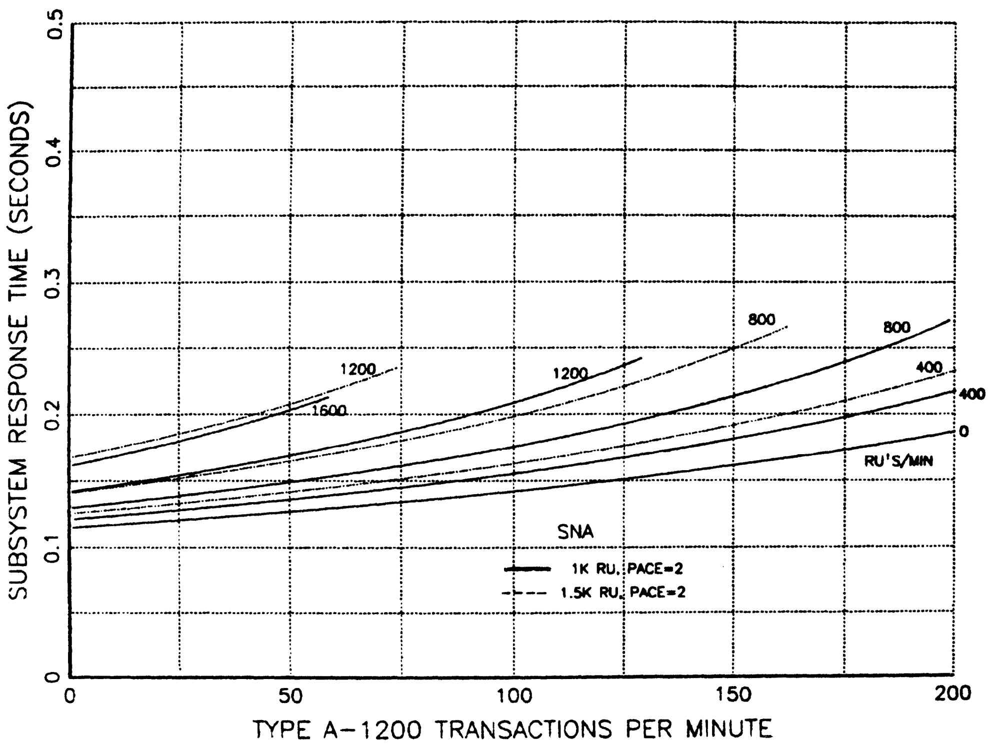

Chart 4-1A: Effect of Long Data Streams on 3174-1L Response, SNA

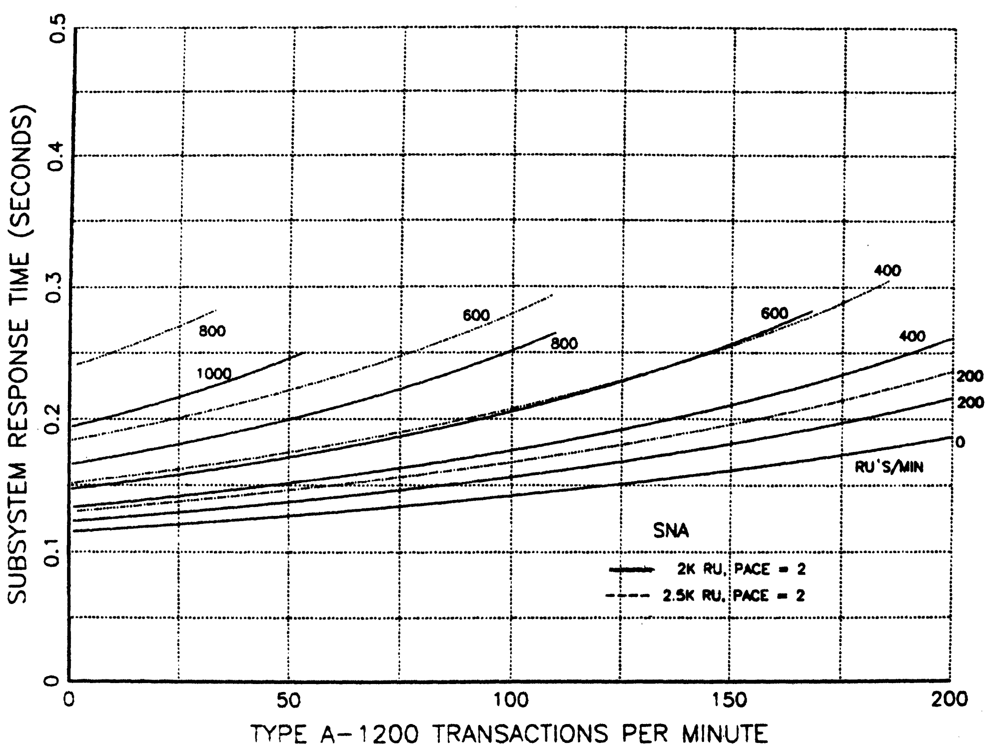

Chart 4-1B: Effect of Long Data Streams on 3174-1L Response, SNA

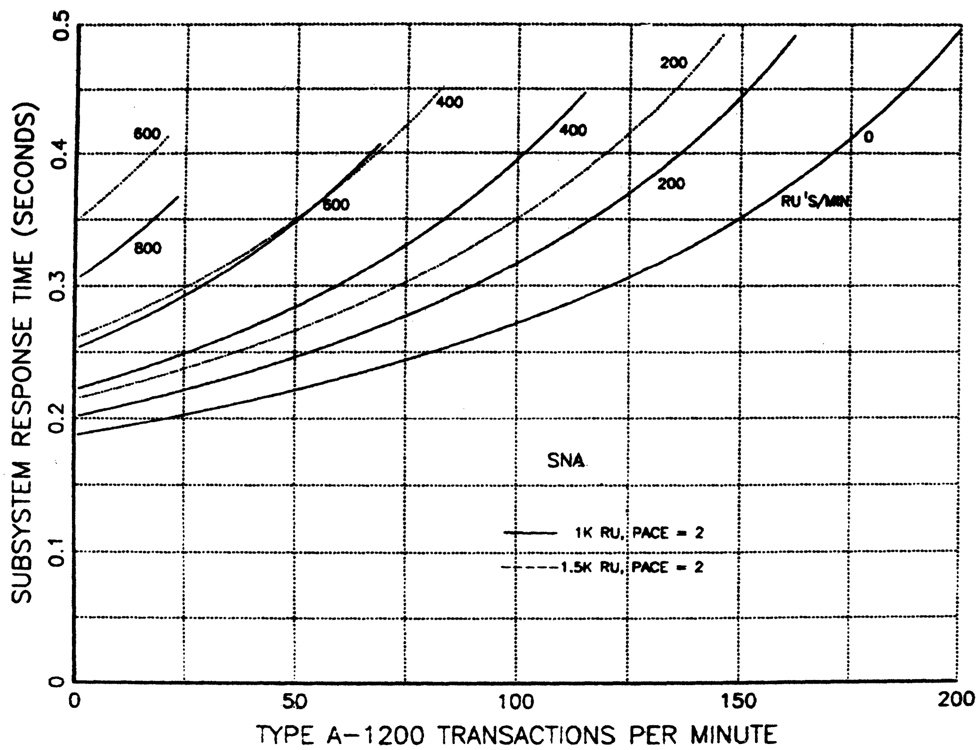

Chart 4-2: Effect of Long Data Streams on 3274-41A Response

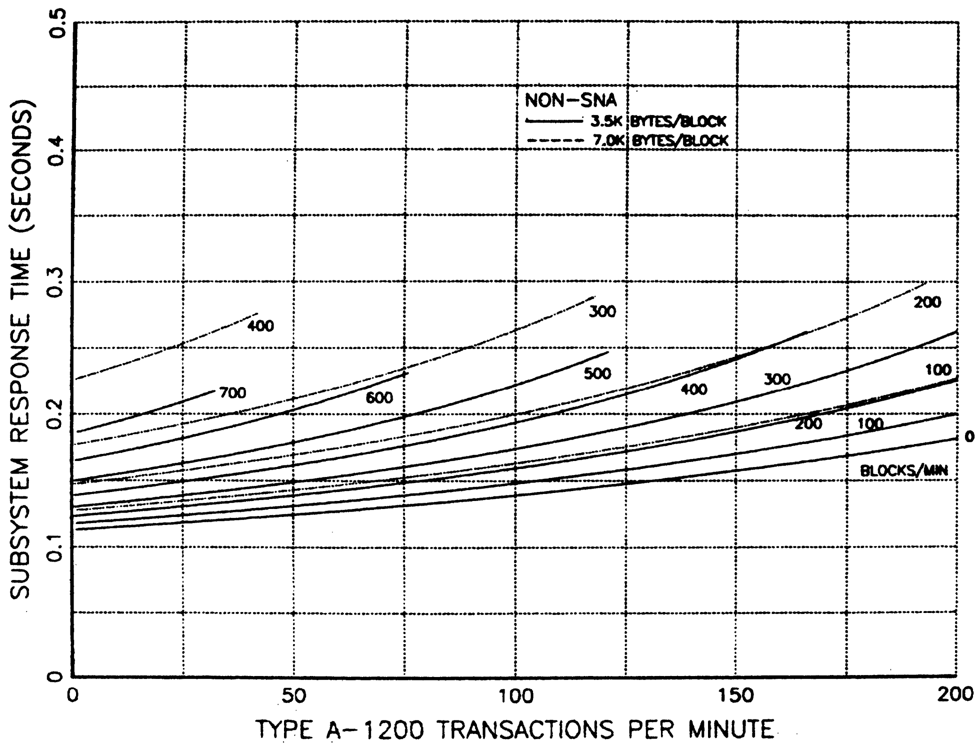

Chart 4-3: Effect of Long Data Streams on 3274-1L Response, non-SNA

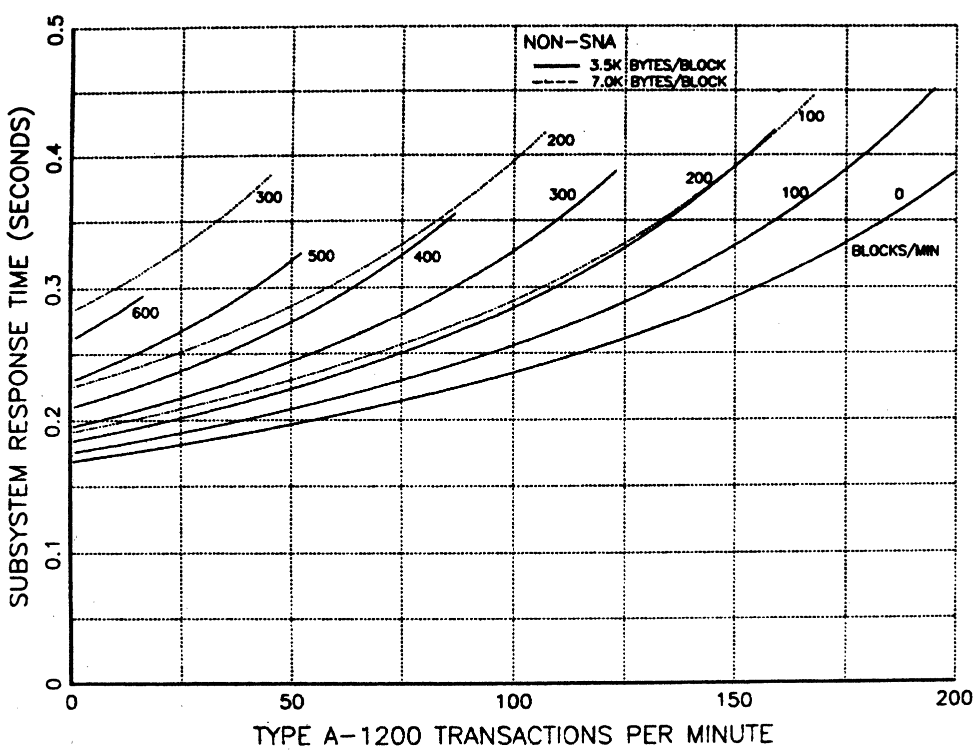

Chart 4-4: Effect of Long Data Streams on 3174-41D Response

CHAPTER 5. DATA STREAM LENGTH AND CONTENT, BENCHMARKS

In addition to length, the content of a data stream affects its service time in an IBM 3X74 subsystem because the execution time of the different commands and orders varies. Thus, for a given subsystem, data stream processing time is a function of the command type (and WCC content, if any), the types and number of orders, and the number of characters being sent.

It is not practical to provide performance data on the basis of data streams associated with specific applications for several reasons:

- The use of display stations is so widespread that it is not feasible to define a reasonable number of 'typical' applications.
- A display station screen is often the joint creation of several programs with different origins.
- Data streams can be written in many different ways to create the same result on the screen.

A number of **alphameric** benchmark data streams have been devised to show a range of performance for some typical combinations of screen size and function. They consist of an identical inbound message and an outbound message to create a 1200, 1560, 2160, 4800, or 5760 character presentation on the display surface (for 24x80, 32x80, 43x80, 90x80, and 62x160 display station buffers, respectively).

The inbound message is initiated by pressing a PF or Enter key, followed by the execution of an RM command bringing in one field of 40 characters.

The outbound message characterizes the benchmark because it tends to be longer, may contain more complex orders, and usually takes more time to process than the inbound message.

The specifics on the composition of these benchmarks are listed in Appendix B. Their intended applications are:

Type A Used for characterizing the performance of monochrome (or four-color) display stations without extended function; contains EW command, SBA and SF orders, and characters.

Type E Used for characterizing the performance of displays with an Extended Attribute Buffer (EAB), for example, seven-color, and/or display stations with extended highlighting. Type E and type A benchmarks are the same, except that all SF orders in type A are replaced by SFE orders.

Type C Used for evaluating control unit ability to handle complex (RA and PT) orders; WRITE command with Reset MDT specified in WCC, with half the SBA and SF orders, and less characters than the comparable type A data stream, but several RA and PT orders.

Type C1 Same as type C but using SFEs and SAs for about twenty-five percent of the fields with EDS (Extended Data Stream) field and character

highlighting attributes added (color attributes used with the 3279 will yield same performance).

The type A and E alphameric benchmarks write on about 83 percent of the rows available on a display station screen, and about 62 percent of the available character positions. These benchmarks aim to represent a typical average transaction, because applications often intersperse many partial screen updates with full screen write operations.

The use of type C and C1 benchmarks assumes that the workstation buffer has been previously loaded with a type A benchmark.

Benchmarks for printers and image displays are provided in Chapter 11 and in the IBM 3193 data sheet, respectively. The benchmark used for file transfer operations is discussed in Chapter 8.

PERFORMANCE-ORIENTED DATA STREAM DESIGN

There is evidence to suggest that not all data streams being used in the field were designed with minimization of subsystem response time in mind. Many take more time than necessary, because they include redundant operations, and/or do not use the most effective means to obtain a desired result.

For local subsystems, both the length and content of a data stream affect service time. Minimization of data stream service time in shared 3X74-based subsystems helps minimize their response time and control unit utilization. Lower utilizations, in turn, help avoid subsystem response time increases that occur as a result of message queues building up.

The benefits of most recommendations in this section will be greatest with CUT mode performance, especially in 3274 control units. In a 3274 operating in CUT mode, certain data stream elements initiate complex transactions over the link that connects the microprocessor in the control unit with the display buffer content in the workstation. These transactions may contribute significantly to response time. In the 3174, these effects are ameliorated because the entire workstation buffer content is retrieved, if necessary, before data stream processing starts, which avoids transactions over the control unit/workstation link during processing. DFT workstations contain both the microprocessor and display buffer, which eliminates the need to use this connection during data stream execution entirely.

For other relevant information, refer to the chapters entitled 'Screen Design' and 'Screen Management' in the **3274 Control Unit Description and Programmers Guide,** GA23-0061.

For remote subsystems, the data stream design criterion is minimization of data stream length irrespective of content, because the transmission time over the telecommunication link usually constitutes the dominant delay. Processing in the subsystem overlaps the transmission interval in most cases. Because minimization of service time and data stream length sometimes lead to contradictory recommendations, the user has the option to differentiate between data streams for local and remote subsystems, to adopt a compromise befitting his requirements and preferences, or to use a data compression or compaction feature, if available. In 3270 data streams, for example, the RA order can be used to compress a series of identical characters.

A paramount design principle is to **avoid specification of unnecessary operations,** for example, avoid obtaining a desired effect more than once. Redundancies are not always obvious; looking for them can be rewarded by significant savings in data stream length and service time. Also, when there are alternative ways to put the same information on a screen, performance can be further improved by selecting the method with the least amount of processing time.

Before discussing the performance aspects of orders and commands, let us examine the design of a display station screen first. Almost all screens are formatted, and are designed to have protected as well as unprotected areas.

The purpose of protected areas is to present information and prompts to the operator, which should not be modifiable by keyboard entries. Protected areas can be implemented by specifying one or more protected fields. Because a field

can extend over several rows, a contiguous area need only be subdivided into separate fields for increasing the intensity or suppressing the display of certain data, or changing selector pen detectability. Because an operator cannot tell the difference between a null and a blank character by looking at the screen, there is no need to differentiate between them in protected fields; see the remarks for the WRITE command below.

Unprotected areas are designed to display keyboard entries, usually one item per input field, and are sized to accommodate them. When a keyboard entry sets the Modified Data Tag (MDT) of a field to "1", its contents (with nulls suppressed) will be transmitted to the host in response to a subsequent Read Modified (RM) command. To permit an operator to use the insert mode, the trailing characters in an input field should be nulls rather than blanks, except for the Entry Assist function.

The choice of orders in outbound data streams has important performance implications, because outbound traffic is usually the major contributor to control unit utilization. For local subsystems, it is recommended to only use the SBA, SF, RA, and IC orders, and the SFE, and SA orders as well when seven-color, extended highlighting, and/or the Programmed Symbol feature are used. For complete information on their operations, refer to one of the 3X74 control unit **Description** manuals.

- **SBA --** The purpose of the Set Buffer Address order is to **change** the current buffer address. Therefore, keep track of how the current buffer address changes during data stream execution and **do not** use an SBA when the current buffer address already has the value you want it to have.

  Because the processing of an SBA order takes approximately as much time as the processing of about sixty 3174 to twenty 3274 characters, insertion of blanks for bridging small gaps in the buffer position sequence will reduce processing time. For large gaps, an SBA order will give better performance.

  For remote subsystems, use of SBAs is indicated for gaps of four or more positions because the emphasis is on minimization of the data stream length.

- **SF --** The processing time of a Start Field order takes approximately sixty 3174 to five 3274 character times, and therefore rewriting an attribute character takes about as much time as skipping over it by using an SBA. Note that rewriting an SF offers the opportunity to reset the MDT. See the remarks for the WRITE command for an example.

- **SFE, MF, SA --** These orders control the content of extended and character attributes.

  The execution time of a Start Field Extended order varies but is always larger than of an SF order; therefore, use an SFE only when an SF cannot accomplish the desired result. Any permissible attribute type not being specifically defined by the SFE order has its value set to binary zero(s).

  The Modify Field order selectively modifies the extension part of extended attributes, and is therefore not used in data streams following Erase Write and Erase Write Alternate commands. In a 3174, for the same number of attribute pairs, the execution times of the MF and SFE orders are about the same. In the 3274, it always takes less time to use an SFE even though all attributes which are not being modified must be restored.

Because character attributes can only be set by the Set Attribute order, there are no performance alternatives to be considered.

- **RA** -- Like an SBA, the Repeat-to-Address order changes the current buffer address but, in addition, inserts a specified character (including null) in all buffer locations being traversed. It can be used to fill part of a row with identical characters, for example, dashes, trailing blanks, or nulls, in an input field.

  Because the processing of an RA order takes in the order of one hundred character times (in both the 3174 and 3274), the trade-off with using a string of identical characters should be carefully considered, especially in local subsystems. In remote subsystems, with their emphasis on limiting data stream length, RAs can be used to compress data streams when they contain strings of five or more identical characters.

  Note that when a new field is intended to be started at the RA-specified address, it is sufficient to only use an SF order; there is no need to use an SBA as well.

- **EUA** -- The Erase Unprotected to Address order inserts nulls in all unprotected buffer locations being traversed to reach the EUA-specified address. Because each EUA order adds almost 700 3174 to 200 3274 character times in CUT mode (insignificant in DFT mode), it is recommended to just write blanks, or use RAs instead. Use this order only in data streams for remote control units where minimization of the number of bytes going over a telecommunication line is important.

- **PT** -- The operation of the Program-Tab order is closely related to EUA-operation, except that its execution ceases at the first data location of the next unprotected field rather than the EUA-specified address.

  Like the EAU order, PTs have relatively high execution times as well (in the order of 300 3174- to 400 3274 character times in CUT mode, insignificant in DFT mode). They are therefore recommended for use in remote control units only.

There are some important performance issues associated with the use of commands as well.

- **EW(A)** -- The Erase Write and Erase Write Alternate commands clear (insert nulls in) all display buffer locations before any new information is written. There is **no** need for specific orders to write nulls in these locations (again).

  In a 3174 CUT subsystem, data stream processing and clearing of the display buffer in the workstation proceed concurrently. With the 3274, data stream processing can start only after the display buffer has been cleared. Depending on workstation type, clearing a buffer takes from 25 to 100 percent of the time to fill it with data.

  When bit 7 in a Write Control Character (WCC) following an EW or EWA command equals "1" (specifying a redundant Reset-MDT operation), it is ignored.

For the best performance of EW data streams in local and remote subsystems, use SBAs and SFs and RAs for writing strings of identical characters, heeding the caveats discussed for these orders.

- **RM** -- The Read Modified command was designed to minimize the inbound data stream length, and therefore causes only the content (with all nulls suppressed) of modified fields (with MDT=1) to be returned to the host.

  The null suppression feature can cause intended spaces to disappear when an operator uses the cursor-move key instead of the space bar. To avoid this problem, you can write blanks in an input field. The effect on performance of using trailing blanks (hex 40) rather than nulls need only be considered for remote subsystems.

- **RB** -- Avoid using the Read Buffer command because it transmits the entire display buffer content inbound, irrespective of how much, if any, modified information it contains. Its use will usually exact a severe performance penalty. It is primarily intended to be used in diagnostic procedures.

> *NOTE:* Users of the XEDIT editor should be aware that its Set Fullread On option uses RB instead of RM commands, and therefore causes the full display buffer content to be transmitted for every inbound transaction, even where the response to an RM command would have been only a few bytes. *This will significantly impair subsystem response.*

- **WRITE** -- The WRITE command will **not** clear the buffer. All data and attributes not being replaced will remain.

  Bit 7=1 in the WCC following a WRITE command will cause a Reset-MDT operation to be executed, irrespective of whether there are MDTs to be reset. It is recommended that this comprehensive operation be used selectively, because in 3274 CUT operations an SF order can reset an MDT in less than two percent of the time taken by a Reset-MDT operation. Execution of a Reset-MDT operation takes relatively less time in DFT mode and 3174 CUT operation.

- **WSF** -- Because the Write Structured Field command encapsulates write-type and other commands, there are no specific performance recommendations associated with its use.

- **Select** -- For optimum performance in local non-SNA 3X74 control units, some commands should be chained to a Select command to minimize channel utilization and obtain optimal subsystem performance.

  Because the selection of the correct Select command is under control of a host program, make sure that the option for this purpose is invoked. (See Appendix C.)

- **EAU** -- The single-byte Erase All Unprotected command is comprehensive, but has a large service time, it inserts nulls in all unprotected fields, resets their MDTs and the AID byte, unlocks the keyboard, and repositions the cursor. It is recommended for use in remote subsystems only.

TRANSACTION DESIGN FOR OPTIMUM PRODUCTIVITY

The previous section discussed ways to keep the service times of data streams in 3270 subsystems to a minimum. Data streams, however, are just the elements that make up a transaction; transactions, in turn, are the steps in the accomplishment of a given task. A broader view is to design transactions in such a way that an operator can complete a task in the shortest possible time, thereby increasing operator productivity.

When presenting information on a screen, for example, the best performance is obtained by using a single command/data stream. Using several data streams instead introduces additional overhead which, if not demanded by the application, should be avoided. An early example of programs violating this principle were those that wrote display screens line-by-line as if they were printers. Eliminating this practice improved performance dramatically.

Another avenue to pursue is reduction of the number of transactions required to complete a task, even though more data may be involved in individual transactions. Lower transaction rates tend to decrease the load of a subsystem (and of a channel, or telecommunication line), and therefore user-perceived response times.

An example of this is a 'heads-down' data entry application where the operator is required to press the Enter key after keying a single field. Using the Tab or New Line key instead allows several fields to be entered before the Enter key is actuated. This drastically reduces the number of interactions with the host, and the likelihood that the operator's 'rhythm' will be disturbed by occasional response delays.

Using the Entry Assist function, now available on almost all 3270 workstations, saves keystrokes and interactions with the host, thereby enhancing an operator's productivity. In addition, it is recommended to associate frequently-used keystroke-saving functions available with many editors, for example, split/join, duplicate, and delete, with PF keys.

CHAPTER 6. PERFORMANCE CHARTS AND DATA SHEET INFORMATION

This chapter addresses the ways in which subsystem response information may be presented. It also provides alphameric performance data on some workstations. Data on other workstations may be found in their performance data sheets.

The response data assume open-loop operation, that is, average transaction rate at a workstation is not affected by system response. However, increases in subsystem (and host) response will tend to increase message interarrival times, and therefore decrease transaction rates. Although this lessens control unit load and response times, there still is a net loss of operator productivity.

Performance data for local non-SNA control units is based on the use of the correct Select command. (See Appendix C.) Where use of a SNA or non-SNA local control unit is not specifically designated, the data applies to SNA operation.

HOW THE PERFORMANCE DATA WAS OBTAINED

Performance data has been obtained either by analytical modeling or by simulation, depending on the complexity of the problem.

In the analytical modeling approach, waiting times are taken into account in internally developed analytical models using the M/M/1 or M/G/1 queueing theorem. According to these models, the 3X74 subsystem services an open (unbounded) queue of transactions in the host system. The models yield response times that include waiting time in the queue, message transfer time into the control unit, and subsystem service time (host and network times not included).

The simulation capability of the RESQ (research queueing) program package has been used to obtain performance data where the model of a subsystem is too complex for analytical solution as, for example, for estimating the effect of file transfer operations on average MFI response times, and vice versa.

PERFORMANCE CHARTS

A performance chart depicts how subsystem **response to last character (RT)** varies with the subsystem control unit transaction rate (CTR) for a specific workstation/ control unit combination. See chart 6-1A for an example. The single-thread service time of a transaction is the RT value at the intersection of a curve with the vertical axis (zero-load condition).

Use of the subsystem RT/CTR curves requires the selection of a representative benchmark, and an estimate of the total control unit transaction load (CTR).

Performance curves have not been extended beyond a control unit utilization of 65 percent because the average delay starts to rise more steeply, and the variability in responses greatly increases. (For curves extending to the edge of a graph, one may assume that the 65-percent point has not been reached.)

RT for a **local** subsystem is defined as the sum of the average processing times of the inbound and outbound messages of an MFI transaction, and of the average queueing delays as a result of control unit utilization. Host processing time is not included, unless indicated.

For CUT mode, service time applies to data stream execution in the control unit, transfer to and from the display station, and transfer over the channel. In DFT mode, service time includes data transfer over the channel and through the control unit, and the data stream processing in the workstation. The latter is independent of control unit utilization, and therefore cluster size.

The **control unit transaction rate (CTR)** (in MFI transactions per minute, for all attached devices) is used as a measure for expressing 3X74 subsystem loading. A given CTR may be produced by a few active terminals with high transaction rates, by many terminals with lower transaction rates, or by a combination. Message destination matters little to the control unit.

The CTR equals the number of actively used terminals (not necessarily all those logged on) multiplied by 60 seconds per minute, divided by the average inter-arrival time (IAT = seconds between transactions). For example, if you have 20 active terminals, each with an average IAT of 10 seconds, then:

CTR = 20 x 60/10 = 120 transaction per minute.

The CTR is likely to vary with the environment. In a laboratory, for example, only 10-20 percent of all terminals on a control unit may be active at any one time, while in a data entry shop 80-90 percent may be active.

EXAMPLE OF HOW TO USE THE CHARTS

Chart 6-1A is based on the use of the type A-1200 benchmark in 3174 and 3274-based subsystems. This data stream approaches full screen output, with a 40 byte input. This is similar to editors (for example, XEDIT, SPF, EDGAR) and simple maps (CICS BMS - given 40 fields).

Suppose you have a 3274-41A with 30 terminals in a production environment. Each operator keys and hits Enter every 10 seconds, on the average. Therefore, CTR = 30 x 60/10 = 180 transactions per minute. From inspection of the chart, using the curve for a 3278-2 attached to a 3274-1L, we see that the 3274 component of response time is about 0.46 seconds.

Changing to a 3174-1L (with a 3278) will drop the subsystem response to about 0.16 second, or about 0.3 second less.

With the same equipment in a laboratory environment, maybe only ten terminals are active, yielding CTR = 60. Here, using a 3174 in lieu of a 3274 would reduce response time by about 0.1 second (from 0.23 to 0.13 seconds).

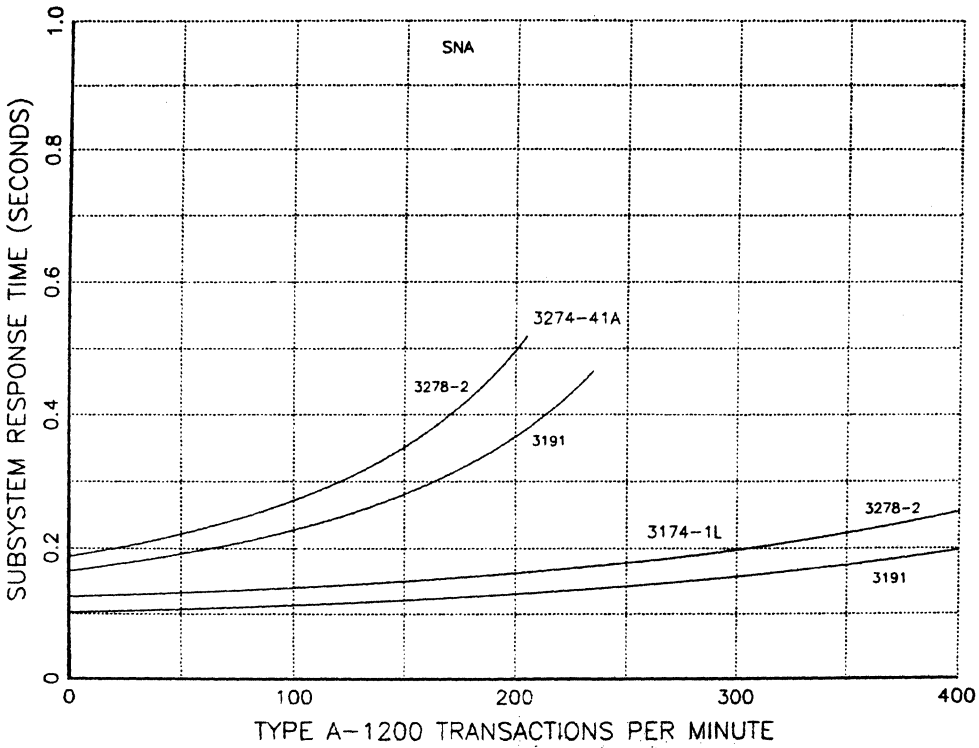

Chart 6-1A: Response of 3X74 with 3191/3278 (SNA, A-1200)

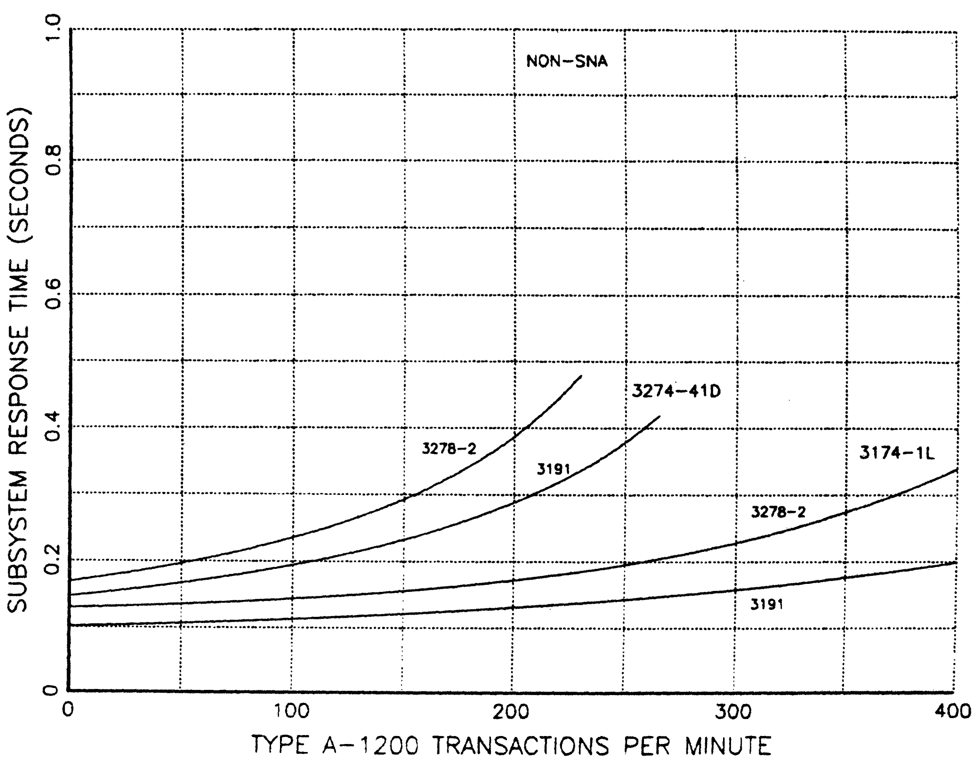

Chart 6-1B: Response of 3X74 with 3191/3278 (non-SNA, A-1200)

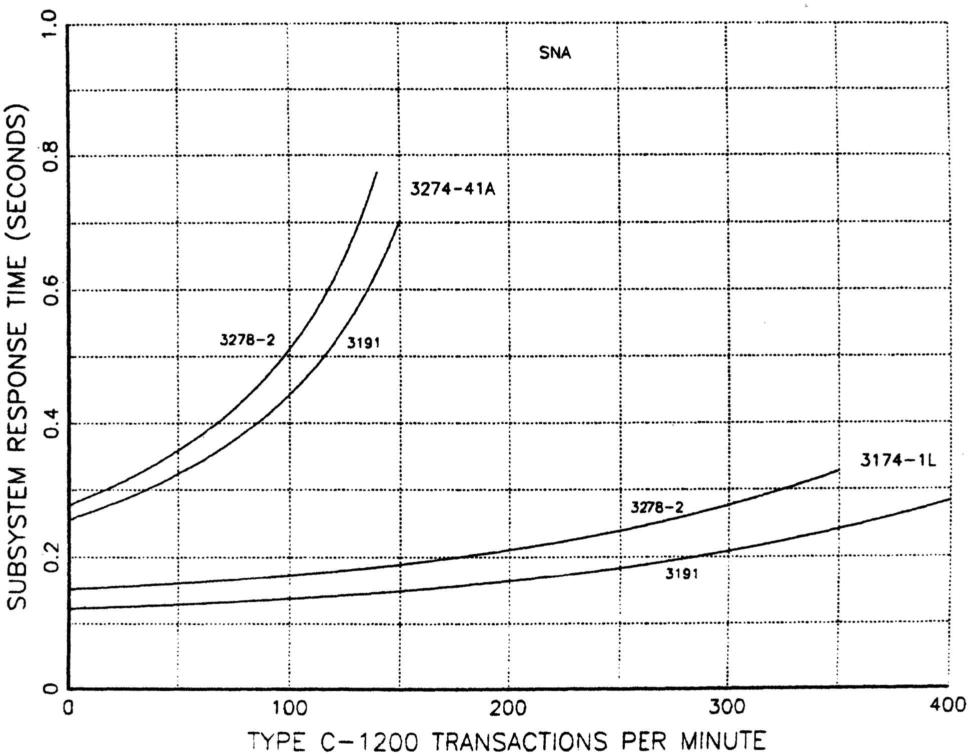

Chart 6-2A: Response of 3X74 with 3191/3278 (SNA, C-1200)

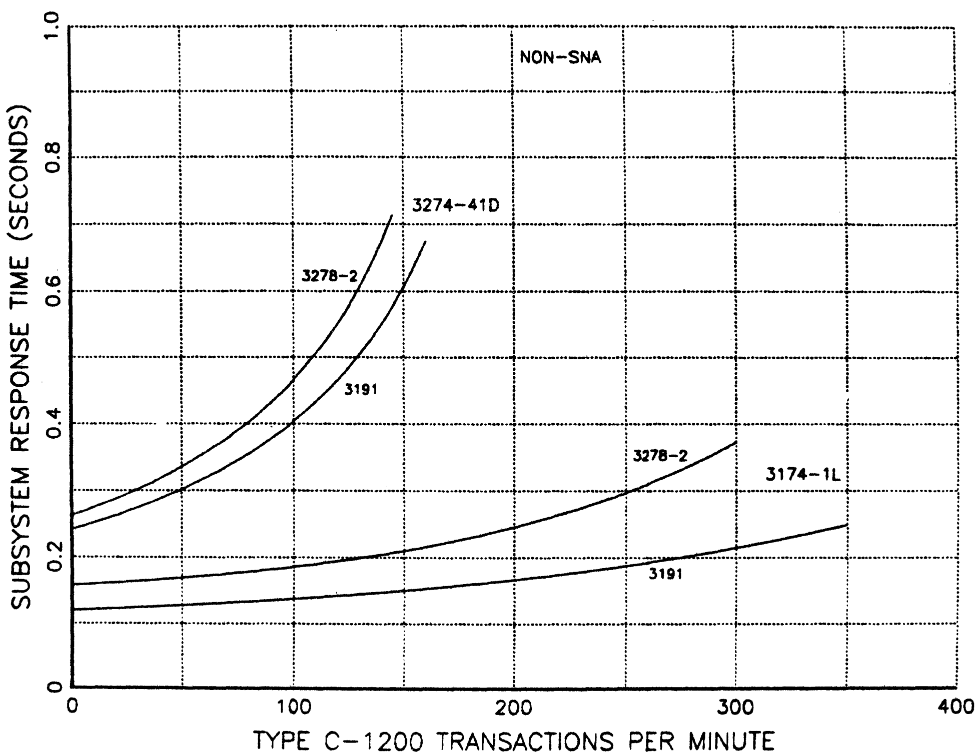

Chart 6-2B: Response of 3X74 with 3191/3278 (non-SNA, C-1200)

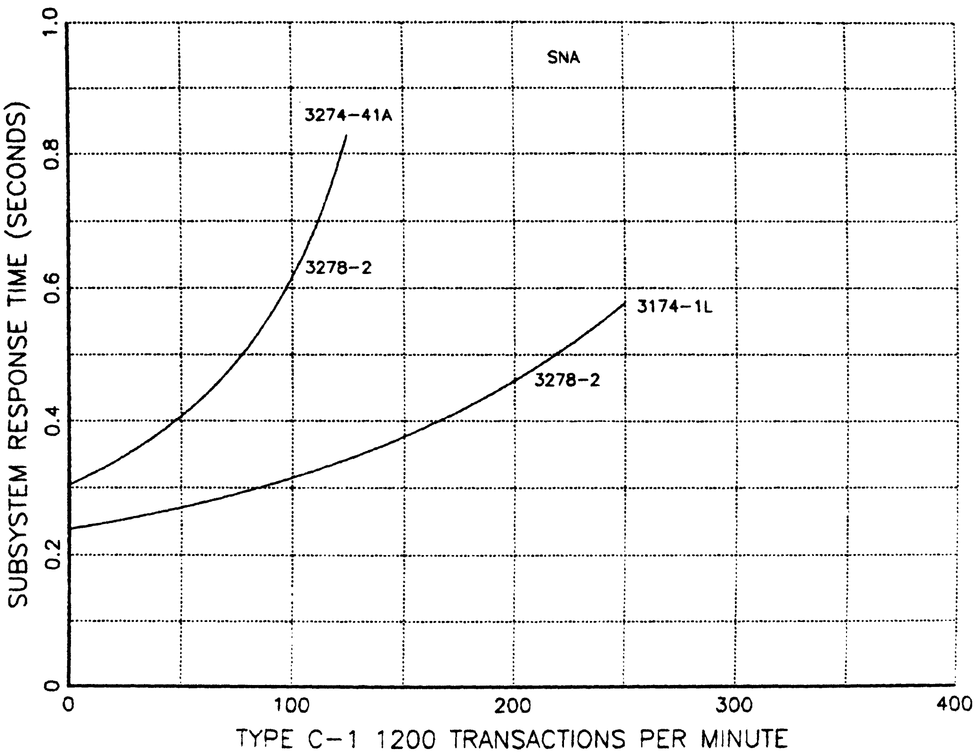

Chart 6-3A: Response of 3X74 with 3191/3278 (SNA, C1-1200)

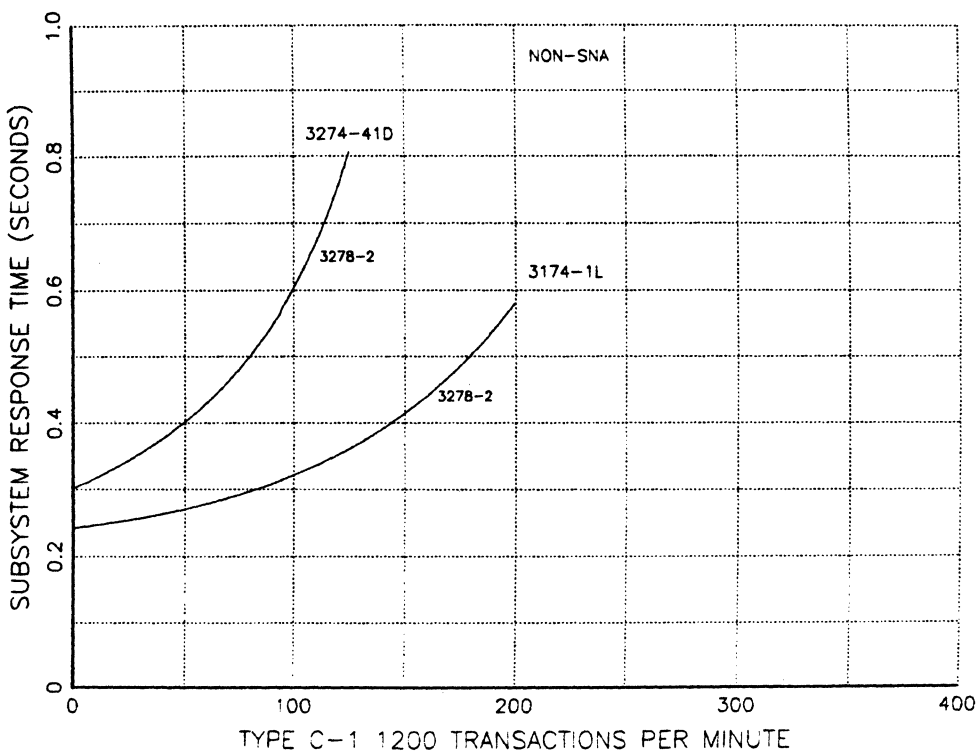

Chart 6-3B: Response of 3X74 with 3191/3278 (non-SNA, C1-1200)

NUMERICAL PRESENTATION OF RESPONSE DATA

Performance information in graphic form as represented by chart 6-1A and others cannot be distributed electronically at this time. The table in Figure 7 provides similar information as contained in charts 6-1A through 6-3B in numerical form, with some differences.

Subsystem response time is shown as a function of control unit utilization rather than CTR. (For 10, 20, 30, 40, 50, 60, and 65 percent average utilization.) This makes the information more versatile because utilization estimates obtained as described in Chapter 7 can be used to estimate subsystem response time for a specific benchmark, rather than a control unit load (CTR) implied by a transaction rate (per minute) of identical benchmarks. Single-thread response times are shown in the 0-percent column.

The table also provides the average control unit utilization (in percent) generated by 6.7 MFI transactions per minute of the specified type, representing the estimated load imposed by one active workstation in a production application.

WS Type	CU Type	Envir't	Benchm.	CU util'n, in %	Subsystem Response Times, in milliseconds at CU utilizations, in %, as follows:							
					00	10	20	30	40	50	60	65
3191	3174	SNA	A-1200	0.94	99	108	120	134	152	178	216	237
3191	3174	non-SNA	A-1200	0.94	101	110	121	136	156	182	227	260
3278	3174	SNA	A-1200	1.01	124	134	147	163	183	214	257	280
3278	3174	non-SNA	A-1200	1.02	127	139	152	170	203	253	333	380
3191	3274-41A		A-1200	1.85	165	183	206	236	275	330	413	471
3191	3274 41D		A-1200	1.64	147	163	184	210	245	294	368	420
3278	3274 41A		A-1200	2.10	187	208	234	267	312	374	468	534
3278	3274 41D		A-1200	1.90	169	188	211	241	282	338	423	483

NOTES:

- "3174" indicates IBM 3174 Model 1L subsystem control unit
- "3274" indicates IBM 3274 Model 41 control unit
- Data does not include contingency
- CU utilization is based on 6.7 transactions per minute.

Figure 7. Response Times of 3X74 with 3191 and 3278 (CUT Mode/A-1200)

HOW TO USE NUMERIC RESPONSE DATA

Assume, for example, that there are 24 3191 active workstations attached to an SNA 3174 Model 1L control unit, each transacting an average of 6.7 A-1200 transactions per minute (CTR = 161). They would cause the 3174 to be about 23 percent (= 24 x 0.94) utilized.

Interpolating between the response times in the "20" and "30" percent column for type A-1200 transactions in 3191/3174 subsystems yields an average subsystem response of about 124 milliseconds.

With an estimated average workload of only four transactions per minute per station, the 3174 utilization would be proportionally less, that is, about 14 percent less, and reduce average response time to about 113 milliseconds.

Using a 3274 Model 41A with 24 3191 workstations (at 6.7 transactions per minute) increases control unit utilization to about 45 percent, and subsystem response to about 0.30 seconds.

CUT MODE PERFORMANCE FOR LOCAL 3X74 CONTROL UNITS

The curves in charts 6-1A through 6-3B, for 3191 and 3278 Model 2 display stations, depict the performance of the 3174 Model 1L and 3274 Model 41 channel-attached control units in both SNA and non-SNA environments.

For Model 2 displays using CUT mode, 3174 response times are significantly better than of the 3274 in both SNA and non-SNA environments. This is illustrated by charts 6-1A, 6-2A, and 6-3A for SNA, and charts 6-1B, 6-2B, and 6-3B for non-SNA.

Note that with the 3174, SNA performance becomes better than non-SNA as transaction rates increase. This is because the internal structure of the 3174 takes advantage of the concurrency inherent within SNA architecture that allows more than one transaction to be processed simultaneously. The 3174 SNA charts assume that the default value for the Attention Delay Value is used (Configuration question 223).

When the file transfer aid configuration bit is set to 1 (digit 6 in question 125, default = 0) to allow PCs with a 3278/3279 emulation adapter to perform file transfer, then the performance of inbound operations on all CUT mode display stations may be degraded. With this bit set to 1, the display buffers must be read by the 3174, even though they may not have been altered. The performance degradation, relative to 3174s with the bit set to 0, may be noticeable to the user for operations (such as paging through a file) in which the display buffer is not altered by the user.

The table in Figure 8 on page 47 lists single-thread response time measurements for some more CUT workstation/benchmark combinations. In many CUT display station types, the last character appears on the screen at virtually the same time as DE or SNA response (+DR) is returned to the host. However, in some display types, the 3180 for example, data transfers between internal buffers in the workstation slightly delay the display of data on the screen with respect to the response to the host, as noted in the table.

The table also lists control unit utilizations for various workstations used for 6.7 transactions per minute per workstation.

For the 3174 Subsystem Control Unit, the results for SNA and non-SNA are approximately the same, while for the 3274 the SNA responses appear to be 10-15 percent more.

IBM 3178/3X74 subsystems have a slightly better response time than the 3278 display station.

IBM 3179/3X74 subsystem response times are about the same as that of 3278 display stations.

Workstation	Benchmark	3174-1L SNA or non-SNA RT	U %	3274-41 SNA RT	U %	non-SNA RT	U %
3191-2	A-1200	97	0.95	165	1.85		
	C-1200	116	1.30	255	2.85		
3178-2	A-1200	111	1.25	193	2.15		
	C-1200	143	1.60	362	4.05		
3179-2	A-1200	116	1.30	187	2.10	169	1.90
	C-1200	150	1.65	276	3.10	262	2.95
	C1-1200	225	2.50	304	3.40	301	3.35
3180[1]	A-1200	103	1.15	163	1.80	137	1.55
	C-1200	123	1.35	259	2.90	241	2.70
	C1-1200	173	1.95	297	3.30	282	3.15
	A-2160	164	1.85	234	2.60	198	2.20
	C-2160	199	2.20	401	4.50	380	4.25
	C1-2160	284	3.15	455	5.10	434	4.85
	A-4800	309	3.45	405	4.50	361	4.05
	C-4800	407	4.55	761	8.50		
	C1-4800	605	6.75	892	9.95		
3278/79-2	A-1200	115	1.30	187	2.10	169	1.90
	C-1200	150	1.65	276	3.10	262	2.95
	C1-1200	228	2.55	304	3.40	301	3.35
	E-1200	172	1.90				
3278/79-3	A-1560	158	1.75				
	C-1560	192	2.15				
	C1-1560	295	3.30				
	E-1560	234	2.60				
3278-4	A-2160	190	2.10	261	2.90	222	2.50
	C-2160	243	2.70	420	4.70	401	4.50
	C1-2160	379	4.25	477	5.35	453	5.05
	E-2160	285	3.20				
3270-AT[1]	A-1200	96	1.05	184	2.10		
	C-1200	115	1.30	267	3.00		

NOTES: RT Single-thread response times in milliseconds

U Utilization in percent, for 6.7 transactions/minute

[1] Does not include time between DE and display on screen

Figure 8. Response Times and CU Utilizations for CUT Workstations

DFT MODE PERFORMANCE FOR LOCAL 3X74 CONTROL UNITS

This section addresses the alphameric performance of local 3174 and 3274 subsystems with display stations operating in DFT mode, and workstations doing 3278/79 emulation in DFT mode.

As discussed in Chapter 2, in DFT operation the workstation rather than the control unit performs keystroke and data stream processing. The control unit merely passes data streams through between the host and a workstation. The performance of **pass-through** in DFT mode was addressed in Chapter 4, including recommendations for long data streams.

The pass-through performance improvement in the 3174 with respect to the 3274 (by 10-25 percent for MFI-type data streams) does not improve the MFI response of many DFT workstations by much. The latter usually process data more slowly than the control unit, and are therefore responsible for a major portion of subsystem response time. However, control unit utilization is reduced, and a heavier transaction rate can be supported.

Chart 6-4* depicts MFI response curves for a local 3274-41 subsystems (using type A benchmarks) that include an IBM 3290 Information Panel. They clearly show the decreased sensitivity to control unit transaction load for this DFT workstation. Attachment to a 3174 will hardly improve these response times. (The asterisk indicates that the data represented in the chart include a 10-15 percent contingency.)

Chart 6-5* shows 3278/79 and 3290 response curves when various mixtures of these display stations are attached to a 3274-41 subsystem (using type A-1200 benchmarks). Note that, for a given transaction rate, response times increase when the 3278 percentage increases because of increased control unit utilization. With the 3174 this effect will be less pronounced, because 3278s use the control unit less.

Charts 6-6A* and 6-6B* depict 3290 and 3278/79 response curves like chart 6-4*, but for type C and C1 benchmarks instead. The relatively better 3290 performance with respect to the 3278 is due to the fact that all processing and buffer storage is confined to the 3290.

The attachment of a workstation operating in DFT mode, like the IBM 3290, requires configuration support A in the 3174 control unit, and configuration support C or D in the 3274.

Figure 9 on page 51 lists control unit utilization (for 6.7 transactions per minute), and single-thread subsystem response times for several **DFT** workstation/benchmark combinations. Note that for this mode of operation, the control unit utilization for a given benchmark appears to be relatively independent of the workstation type. (For example, from 0.5 to 0.6 percent for 6.7 type A-1200 transactions per minute in a 3174.)

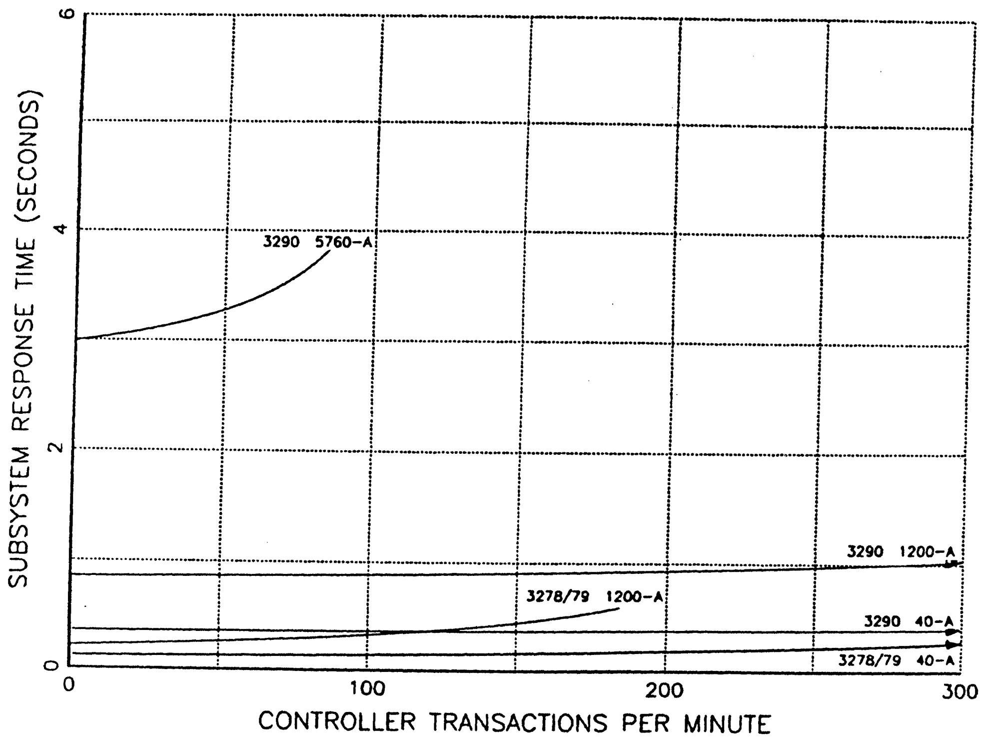

Chart 6-4*: Response of 3290 and 3278 on 3274-41, Type A DS

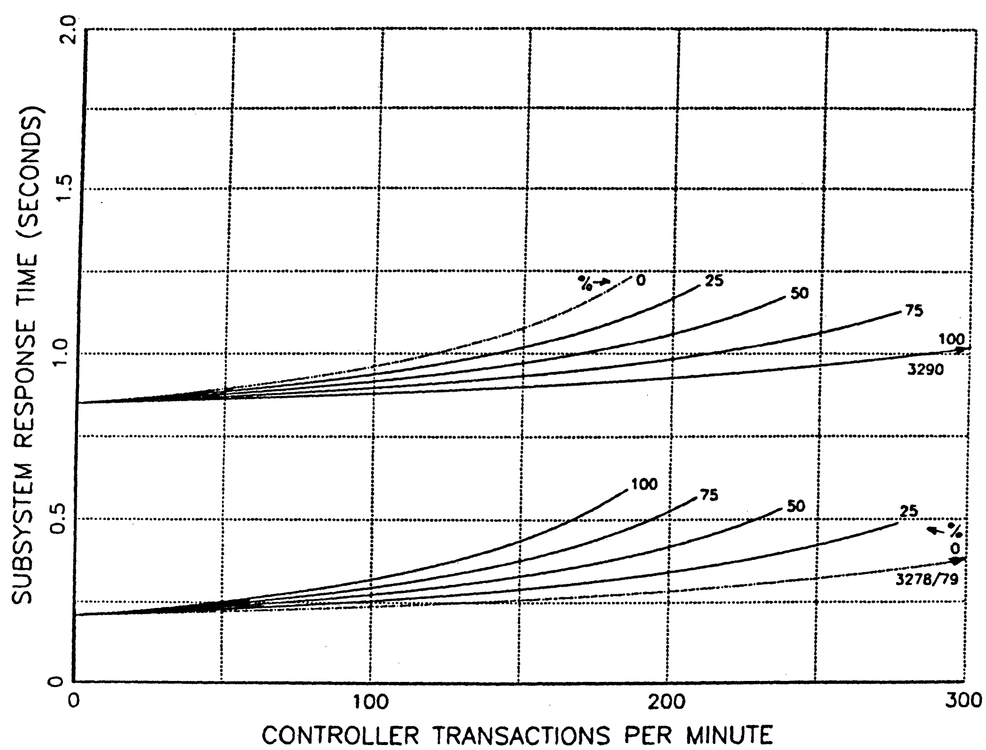

Chart 6-5*: Response of 3290/3278 Mix on 3274-41

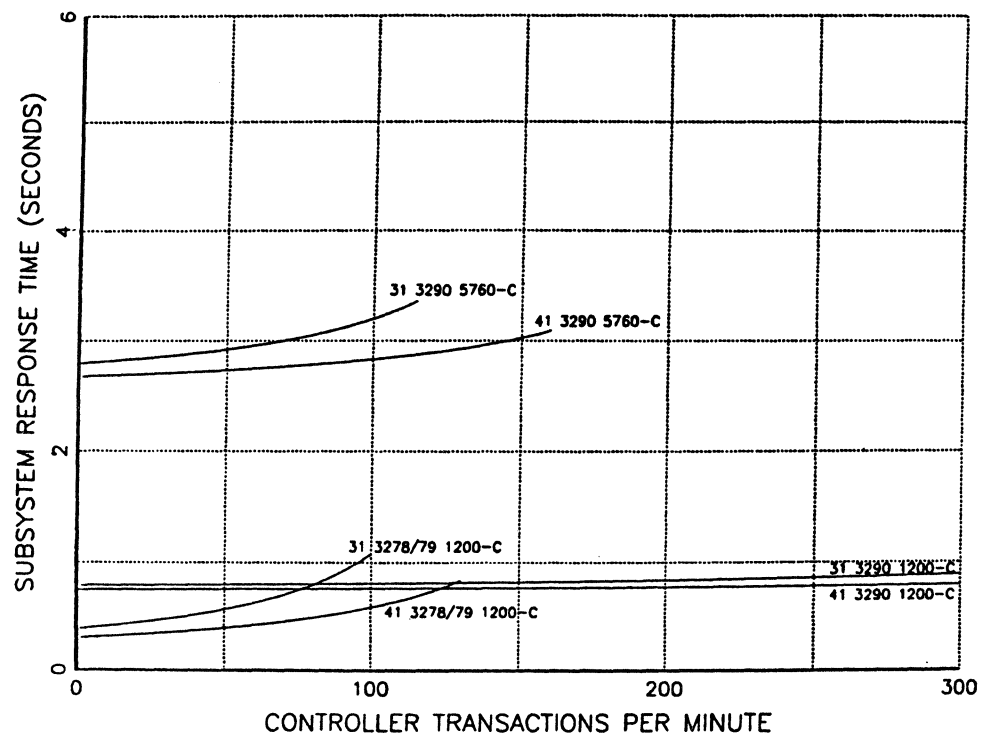

Chart 6-6A*: Response of 3290 and 3278 on 3274-41, Type C DS

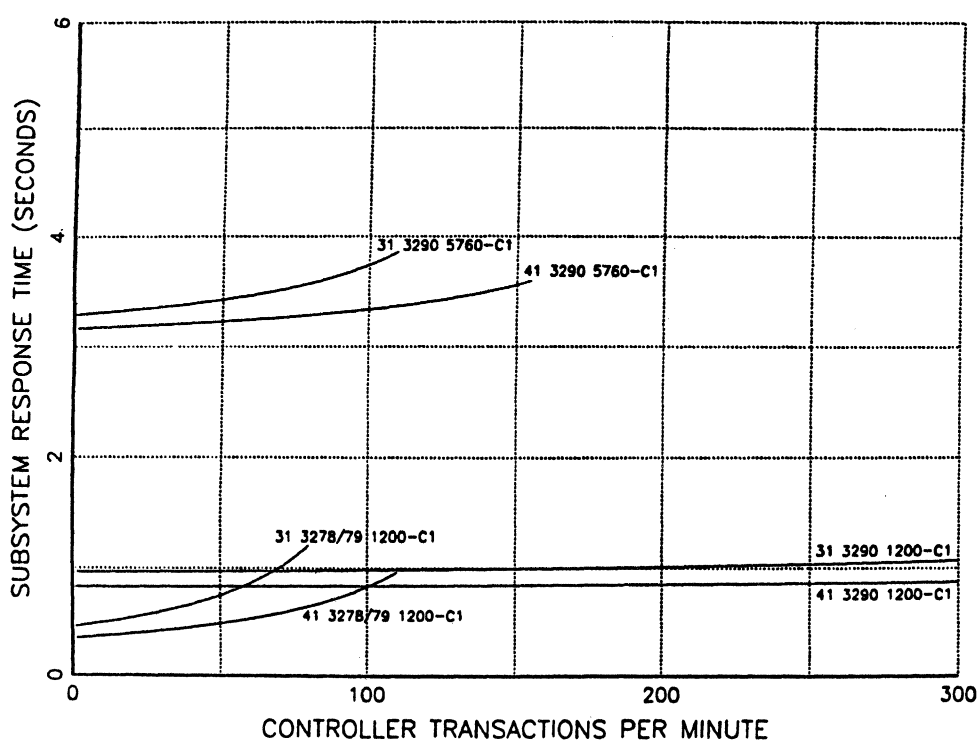

Chart 6-6B*: Response of 3290 and 3278 on 3274-41, Type C1 DS

Work-station	Bench-mark	3174-1L SNA RT	3174-1L SNA U %	3174-1L non-SNA RT	3174-1L non-SNA U %	3274-41 SNA RT	3274-41 SNA U %	3274-41 non-SNA RT	3274-41 non-SNA U %
3290-	A-1200	454	0.60	530	0.79	482	0.89		0.90
	C-1200	333	0.54	431	0.74	351	0.67		0.83
	C1-1200	353	0.53	457	0.74	374	0.68		0.83
	A-5760	1573	1.57	1530	1.28				
	C-5760	1173	1.02	1242	1.01				
	C1-5760	1296	1.02	1339	1.03				
3179-G	A-1200	154	0.50	206	0.77				0.87
	C-1200	259	0.42	320	0.74				0.83
	C1-1200	262	0.42	321	0.75				0.82
3193	A-1200	415	0.60	439	0.77	441	0.90	456	0.86
	C-1200	220	0.51	248	0.74	236	0.69	263	0.83
	C1-1200	224	0.52	252	0.75	241	0.70	270	0.83
3270-AT	A-1200	197	0.59	262	0.78	254	0.93	277	0.84
	C-1200	194	0.52	266	0.75	238	0.69	280	0.75
	C1-1200	201	0.53	271	0.74	241	0.71	287	0.77
E3270-AT	A-1200					403	0.92	420	0.83
	C-1200					341	0.70	366	0.78
3270XT/GX	A-1200	463	0.60	524	0.78	486	0.92		
	C-1200	439	0.51	509	0.74	451	0.70		
	C1-1200	451	0.52			464	0.70		
3270AT/G	A-1200	221	0.59	263	0.79	250	0.92		
	C-1200	213	0.52	259	0.73	226	0.69		
	C1-1200	215	0.52	263	0.75	231	0.70		

NOTES:

RT Single-thread response times in milliseconds, with -1200, -1560, and -2160 for model 2, 3, and 4 screens respectively; does not include time between response to host and display of last character on screen.

U Utilization in percent, for 6.7 transactions/minute

Figure 9. Response Times and CU Utilizations for DFT Workstations

CHAPTER 7. LOCAL 3X74 SUBSYSTEM CAPACITY PLANNING

The processing of an inbound or outbound message, or the pass-through of a printer data stream or a file, all represent **work** for a **local** 3X74 control unit, and is quantified in this document as "utilization". This chapter will show how the utilization of a subsystem control unit may be used to assess the capacity and performance of local subsystem configurations.

Note that for **remote** subsystems, the speed, type, and utilization of the communication link determine subsystem performance because control unit processing keeps up with and overlaps the data flow over the link. (See Chapter 9.)

The 3X74 utilization of an operation is the total of all 3X74 service times, in seconds, associated with that operation per second, multiplied by 100 to obtain the utilization in percent. For example, if a workstation processes 6.7 transactions per minute, each requiring 108 milliseconds of 3X74 time to be serviced, then Uws, the 3X74 utilization by that workstation, would be:

$$Uws = ((6.7 / 60) \times 0.108) \times 100 = 1.2 \text{ percent}$$

The capacity planning approach for local 3X74 subsystems is based on adding the contributions of the various MFI transactions, concurrent file transfers, and printer operations to obtain total 3X74 subsystem utilization. This utilization is assumed to be the sum of the utilizations contributed by all operations. In an SNA control unit, the servicing of some parts of more than one data stream may overlap because some of its elements may sometimes operate concurrently. Therefore, the actual utilization of such a control unit may sometimes be slightly less than the sum of the utilizations obtained from single-thread service times. This effect will be ignored in our considerations.

With most CUT workstations, control unit service time approaches subsystem response time but never exceeds it. With DFT workstations, control unit service (pass-through) time is only part of the subsystem response time because during the time that the workstation processes its data stream, the control unit is free to do other things.

When control unit utilization increases, the tendency for queues to form increases. Because the servicing of work in a queue is delayed, the average subsystem response time will increase. Also, because the interference of messages with each other will increase, variability of response times will increase.

The average control unit utilization is a measure for the aggregate load that a local subsystem can handle with acceptable response times. As a general rule, it is recommended that utilization does not exceed fifty percent, or be less than twenty-five percent for applications where minimal subsystem response time is important. For workstations operating in CUT mode, keystroke processing adds to 3X74 utilization; in DFT mode operation it does not.

As was noted in Chapter 1, the contribution of data transfers over the channel to response time is small. High channel utilization, however, can introduce significant delays. It is therefore recommended to keep the average utilization

of a channel by all its attached subsystems and devices below thirty percent. (See Appendix C.)

CAPACITY PLANNING DATA

The control unit utilization percentages, provided in Figure 8 on page 47 and Figure 9 on page 51 in for assisting with 3X74 capacity planning, are based on the assumption that a workstation operator enters an average of 6.7 MFI A-1200 transactions per minute (about 400 per hour). If this too high for your application, you may adjust these percentages. For example, when only 300 transactions per hour, or an average of five per minute is anticipated, and for a given subsystem/benchmark combination a 1.2 percent utilization is specified, use only 0.9 percent per workstation.

These tables also include the associated single-thread subsystem response time measurements, that is, with no other activity in the control unit. For the 3174 control unit, the results for SNA and non-SNA are approximately the same, while for the 3274 the SNA response times appear to be 10-15 percent more.

Note that for DFT mode of operations, the control unit utilization for a given benchmark appears to be relatively independent of the workstation type. (For example, between 0.5 and 0.6 percent for 6.7 type A-1200 transactions per minute in a 3174 subsystem control unit.)

For estimating utilization percentages associated with file transfers and the operation of printers, consult the applicable sections in Chapters 8 and 11, respectively.

As noted before, the use of 2- and 3-digit numbers in these examples does not imply that the obtained estimates are accurate to this degree. They are approximations for the purpose of providing some insight into 3X74 subsystem capabilities.

CAPACITY PLANNING EXAMPLES

To illustrate the procedure, two local subsystems and their transaction loads have been postulated in Figure 10 on page 55.

- The heading "Load type" specifies the operation as CUT or DFT, and whether the load is viewed as open-loop or closed-loop.
- The heading "Workstation" specifies the types and number of attached devices.
- The "Transaction" column specifies the type of benchmark used, and the number of active devices (of those attached).
- The sections for the 3174 and 3274 control units list control unit utilization Uws per active workstation of the specified type, used at a rate of 6.7 transactions per minute. The column on the right contains control

unit usage per line entry, and their total, that is, the aggregate control unit utilization.

- The row "3191 response" at the bottom of each example indicates the 3191/3X74 responses at a control unit load of 0, and the estimated total.

Example	Load type	SNA Environment: Workstation Type	WS	SNA Environment: Transaction Type	WSa	3174-1L Utilization Uws %	3174-1L U %	3274-41A Utilization Uws %	3274-41A U %
#1	CUT/open-loop	3191	20	A-1200	16	0.95	15.2	1.85	29.6
		3278-4	10	A-2160	8	2.10	16.8	2.90	23.2
				Total			32.0		52.8
	3191 response (0-load/U-load), in ms					99 / 137		165 / 350	
#2	CUT/open-loop	3191	15	A-1200	10	0.94	9.4	1.85	18.5
		3278-4	10	A-2160	5	2.10	10.5	2.90	14.5
	DFT/open-loop	3270-AT	5	A-1200	2	0.60	1.2	0.90	1.8
				Open-1p totl			21.1		34.8
	DFT/clsd-loop	3270-AT		FT	2		33		42
				Total			54		77
	3191 response (0-load/U-load), in ms					99 / 193		165/>500	

NOTES:
- ms — milliseconds
- Uws — Utilization in percent, for 6.7 transactions per minute, per workstation. (Numbers from tables in Chapter 6)
- U — Share/total utilization, in percent
- WS — Number of attached workstations of this type
- WSa — Number of active workstations of this type
- FT — Number of concurrent file transfers

Figure 10. 3X74 Subsystem Capacity Planning Examples

Commenting on the two tabulated examples:

- **Example 1:** The 3174 is estimated to have less than half the 3274 subsystem response time, although either control unit may be used. Furthermore, its lower utilization will make the 3174 subsystem less sensitive to further load increases.

- **Example 2:** Estimation of the open-loop load follows the pattern of Example 1. For estimating the utilization as a result of the two concurrent

file transfer, consult Chapter 8. (Figure 12 on page 61 and Figure 13 on page 62.)

From the performance data sheet for IBM Personal Computer workstations, we obtain 0.20 and 0.30 (approximate) for the AT Bc/B ratios (Chapter 8) associated with the 3174 and 3274 control units, respectively.

For the 3174, in the "20 percent open-loop utilization" column of Figure 13 on page 62 for Bc/B=0.200, we find 33 percent additional utilization for two concurrent file transfers. The E-table in Figure 12 on page 61 indicates an expansion factor of about 1.11, that is, a slow down from the minimum rate of about eleven percent.

For the 3274, from the "30/40 percent" columns in Figure 13 on page 62and the line for Bc/B=0.300, we find 42 percent additional utilization for two concurrent file transfers, and for E about 1.35 (from Figure 12 on page 61), with these conditions.

The total average utilizations add up to about 54 and 77 percent for the 3174 and the 3274, respectively. Note that the MFI responses of the 3274 can expected to be considerably worse, while the file transfer rates will be lower too.

CHAPTER 8. UP- AND DOWNLOADING OF FILES (FILE TRANSFER)

Many PC- and PS/2-based workstations can be attached to IBM 3X74 subsystem control units by using the 3278/79 Emulation adapter (#5050) and 3270 Connection (#2000), respectively. Because such workstations have their own file devices, they impose additional performance demands on the control unit arising from the up- and downloading of files and documents. (Or, demands on the telecommunications link for remote subsystems: however, that is not addressed in this document.)

Performance aspects of these file transfer operations are:

- The time required to initiate the file transfer operation
- The actual time for transferring the file data (transfer time)
- The time required to terminate the operation
- How other transactions (including other file transfers) in the control unit (or host link for remote subsystems) affect transfer time
- The effect of one or more simultaneous transfers in a subsystem on the interactive responses of other workstations in the cluster.

File transfer operations are supported by a program in the host, and a matching program in the workstation that may include CUT and/or DFT mode emulation capability, as available or preferred. Most of these programs divide a file into data blocks of about 2 Kb or 3.5 Kb, and use existing 3270 data stream commands and orders to transfer the file and associated control data in either direction. Therefore, no special provisions are needed in the control unit.

Actual file transfer rates depend on many factors, such as: workstation and control unit hardware, program support in host and workstation, file transfer protocol, and host and control unit loading.

The documentation for various program products may refer to up- and downloading of files in different terms, such as: send and receive, write and read, and inbound and outbound. To allay confusion, the latter terms are often used when discussing file transfer operations associated with a particular product.

FILE TRANSFER OPERATION (LOCAL CONTROL UNITS)

The roles of host, workstation, and control unit during file data transfer are depicted in Figure 11 on page 58; only one of these is involved at any one time (ignoring small overlaps). File transfer messages are interleaved with messages of other workstations on a first come - first served basis.

In our consideration of file transfer, the average service times plus delays in the host and the workstation are assumed to be **independent** of the traffic flow through the control unit. They should be chosen in concert with the type and load of the host and workstation, respectively. Minimal host times will

be experienced with lightly loaded, fast host systems, and program support with short path lengths. Conditions for minimal workstation times are a fast microprocessor, high priority and short path lengths for the code supporting communication with the control unit, adequate RAM buffer area, and a lightly loaded file device.

The delay in the control unit, on the other hand, is considered to be **dependent** on its utilization by other transactions such as MFI, printing, and file transfers.

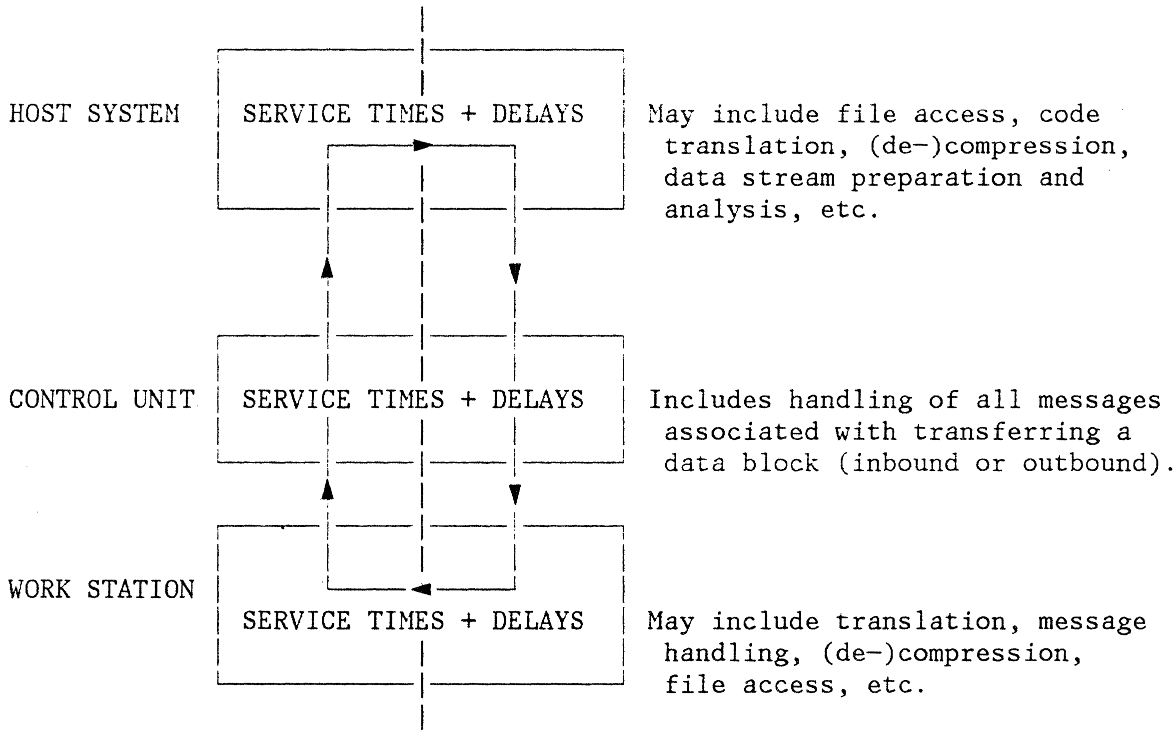

Figure 11. File Transfer Overview

A closed-loop model is used for obtaining the file transfer performance data in these guidelines. Although file transfer implementations functionally emulate operator-initiated interactivity, their performance characteristics differ. A workstation responds automatically, usually within a short, fixed time, to host-initiated messages. An operator may take many seconds to react.

That part of the control unit load contributed by activities other than "closed-loop" file transfer operations, such as MFI transactions and printing, will be referred to as its "open-loop" load. Its utilization of the control unit is considered to be independent of total control unit utilization, although in reality this cannot hold for high total utilization numbers. MFI transaction rates and printer throughput are bound to decrease at some point.

Service times and delays in connection with file transfer operation are expressed on the basis of 'kbyte of file data' transferred.

MINIMUM FILE TRANSFER TIME

The minimum time "Tfmin" to transfer a file in a given environment (host, workstation) is achieved when there is no other activity in the control unit:

$$Tfmin = A + (B \times F) = A + (Bh + Bc + Bw) \times F \text{ seconds}$$

in which:

Tfmin	minimum total elapsed file transfer time, in seconds
A	initiation plus termination time, in seconds
B	transfer time per kbyte, equals Bh + Bc + Bw, seconds/kbyte
F	file length, in kbytes
Bh	non-overlapped time in host per kbyte, seconds/kbyte
Bc	non-overlapped service time in CU per kbyte, seconds/kbyte
Bw	non-overlapped time in workstation per kbyte, seconds/kbyte

and for use later on:

Bc'	non-overlapped total time in CU per kbyte, in seconds/kbyte
Tf	total elapsed file transfer time, in seconds
FTR	file transfer rate , in kbytes/second (= 1/B)

Note that the constant "A" (in seconds) accounts for the total time to initiate and terminate the file transfer operation. (In rare cases, this may be somewhat affected by file size.) Both "A" and "B" vary with control unit model, local or remote attachment, SNA or non-SNA, CUT or DFT mode, whether the data is up- or downloaded, workstation hardware, file device type, the file transfer program support in host and workstation, and speed and load of the host.

EFFECT OF THE SUBSYSTEM CONTROL UNIT LOADING

For given host (Bh) and workstation (Bw) times, and an expanded average control unit time Bc' that includes both service time and delay, a formula for "Tf" based on the expression for "Tfmin" can be written as follows:

$$Tf = A + (B \times E) \times F = A + (Bh + Bc' + Bw) \times F \text{ seconds}$$

This equation defines the expansion factor "E", that is, the factor by which the minimum elapsed time of the transfer phase needs to be multiplied to obtain an estimate for the actual time.

As noted, "Bc'" represents the sum of the control unit service time "Bc" and the average delay as a result of its transaction load; the contributions "Bh" and "Bw" to "B" are treated as constants.

This expression, and the relationship "Bc' = Bc/(1-Uc)" from the M/M/1 queuing theorem, yields:

$$B \times E = Bh + Bc' + Bw = Bh + \frac{Bc}{1 - Uc} + Bw$$

in which "Uc" (expressed as a fraction rather than a percentage) is the total 3X74 utilization by file transfer and open-loop traffic (MFI, printers, etc.). The equation can be rearranged to:

$$E = 1 + \frac{Bc}{B} \times \frac{Uc}{1 - Uc}$$

in which "Bc/B" is the ratio of the service time "Bc" per kbyte of file data transferred in an unloaded control unit, and "B = Bh + Bc + Bw".

This formula is the basis for the tabulation in Figure 12 on page 61. For a given value of "Bc/B", the table provides "E" for 0, 10, 20, 30, 40, 50, 60, and 65 percent open-loop utilization, for 1, 2, 3, and 4 simultaneous file transfers.

The associated values of the total control unit utilization Uc are tabulated in Figure 13 on page 62. Uc is the sum of the utilization as a result of the file transfer(s), and the open-loop activities in the control unit as indicated in the column heading. The utilization as a result of the file transfers alone may be found by subtracting the latter number from Uc.

To understand the effect of one or more file transfer operations on the performance of other transactions, consider this. For each workstation transferring a file, there is never more than one message en route, while for n concurrent transfers by n workstations, the control unit will never need to handle more than n file transfer messages at any one time. Therefore, file transfer operations can delay, but will not totally inhibit control unit operations in support of other devices, although control unit utilization, even for a single file transfer, can sometimes be very high.

Furthermore, each additional file transfer will increase the duration of all transfers in progress which, in turn, increases the probability of even more transfers to run concurrently. Therefore, increasing the combined file transfer load will create a progressively poorer environment for interactive users. As a result, response times will go up and their transaction rate (and control unit utilization) decrease. This situation should be avoided by appropriate subsystem configuration.

The reciprocal of "B x E" is the adjusted file transfer rate (FTR') in kbytes per second, during the transfer phase of the operation:

$$FTR' = \frac{1}{B \times E} \text{ kbytes per second}$$

File Transfer Expansion Factor "E"									
	# of FTs	CU utilization by other CU activity, in percent							
Bc/B		00	10	20	30	40	50	60	65
0.000	any	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
0.100	1	1.000	1.005	1.026	1.057	1.078	1.117	1.163	1.203
	2	1.010	1.019	1.041	1.073	1.098	1.144	1.221	1.250
	3	1.021	1.036	1.060	1.093	1.128	1.181	1.284	1.319
	4	1.033	1.056	1.083	1.118	1.168	1.238	1.365	1.425
0.140	1	1.000	1.014	1.026	1.060	1.099	1.144	1.225	1.253
	2	1.020	1.039	1.056	1.098	1.144	1.212	1.315	1.361
	3	1.043	1.068	1.095	1.143	1.199	1.285	1.431	1.506
	4	1.069	1.101	1.143	1.195	1.279	1.400	1.590	1.701
0.200	1	1.000	1.027	1.057	1.086	1.139	1.210	1.310	1.360
	2	1.034	1.065	1.112	1.162	1.233	1.327	1.485	1.578
	3	1.089	1.124	1.179	1.248	1.349	1.496	1.729	1.838
	4	1.146	1.203	1.276	1.373	1.505	1.717	2.042	2.129
0.300	1	1.000	1.038	1.070	1.146	1.208	1.314	1.491	1.581
	2	1.079	1.137	1.199	1.283	1.406	1.569	1.821	1.969
	3	1.198	1.274	1.369	1.490	1.667	1.905	2.301	2.518
	4	1.350	1.455	1.597	1.766	2.006	2.350	2.906	3.195
0.400	1	1.000	1.049	1.115	1.185	1.271	1.420	1.621	1.725
	2	1.146	1.221	1.317	1.442	1.594	1.831	2.194	2.362
	3	1.364	1.477	1.618	1.796	2.036	2.395	2.942	3.279
	4	1.636	1.794	1.979	2.280	2.583	3.063	3.844	4.256
0.600	1	1.000	1.072	1.161	1.277	1.416	1.633	1.915	2.102
	2	1.343	1.463	1.622	1.829	2.096	2.450	3.046	3.330
	3	1.806	1.999	2.236	2.533	2.954	3.506	4.372	4.862
	4	2.445	2.597	2.930	3.325	3.860	4.642	5.846	6.404
0.800	1	1.000	1.108	1.222	1.367	1.578	1.839	2.235	2.453
	2	1.636	1.805	2.030	2.297	2.676	3.186	3.997	4.451
	3	2.395	2.649	2.970	3.388	3.938	4.717	5.935	6.575
	4	3.167	3.515	3.957	4.497	5.256	6.297	7.937	8.749
NOTE:	FTs	Number of concurrent file transfers							

Figure 12. File Transfer Expansion Factor "E"

FILE TRANSFER IN CUT OR DFT MODE?

You should be aware that one kilobyte of data, as stored on a file device, does not necessarily mean that a file transfer operation transfers the identical number of bytes through the control unit.

"Total" and "File Transfer" Control Unit Utilizations "Uc" and "Uft"

Bc/B	# of FTs	00	10	FT	20	FT	30	FT	40	FT	50	FT	60	FT	65	FT
		"Open-loop" and "File Transfer" Utilizations, in percent														
0.100	1	10	20	10	30	10	39	9	49	9	59	9	68	8	73	8
	2	18	28	18	37	17	47	17	56	16	66	16	75	15	79	14
	3	25	35	25	44	24	53	23	62	22	71	21	80	20	84	19
	4	31	41	31	50	30	59	29	68	28	77	27	85	25	88	23
0.140	1	14	24	14	33	13	43	13	53	13	62	12	71	11	75	10
	2	25	34	24	43	23	53	23	62	22	71	21	79	19	82	17
	3	34	43	33	52	32	61	31	70	30	78	28	86	26	89	24
	4	43	52	42	61	41	69	39	77	37	84	34	91	31	93	28
0.200	1	20	30	20	39	19	48	18	57	17	66	16	75	15	79	14
	2	35	44	34	53	33	61	31	69	29	77	27	85	25	87	22
	3	47	56	46	64	44	72	42	79	39	85	35	91	31	93	28
	4	59	66	56	73	53	80	50	86	46	91	41	95	35	97	32
0.300	1	30	39	29	48	28	56	26	64	24	72	22	80	20	84	19
	2	51	59	49	67	47	74	44	80	40	86	36	91	31	93	28
	3	66	73	63	79	59	84	54	89	49	93	43	96	36	97	32
	4	79	84	74	88	68	92	62	95	55	97	47	99	39	99	34
0.400	1	40	48	38	56	36	64	34	71	31	78	28	84	24	87	22
	2	65	71	61	77	57	82	52	87	47	91	41	95	35	96	31
	3	81	85	75	89	69	92	62	95	55	97	47	98	38	99	34
	4	90	93	83	95	75	97	67	98	58	99	49	99	39	99	34
0.600	1	60	66	56	72	52	78	48	83	43	87	37	91	31	92	27
	2	85	88	78	91	71	94	64	96	56	97	47	98	38	98	33
	3	95	96	86	97	77	98	68	99	59	99	49	99	39	99	34
	4	98	99	89	99	79	99	69	100	60	100	50	100	40	100	35
0.800	1	80	84	74	87	67	90	60	92	52	94	44	96	36	97	32
	2	96	97	87	98	78	99	69	99	59	99	49	99	39	99	34
	3	98	99	89	99	79	99	69	99	59	99	49	99	39	99	34
	4	99	100	90	100	80	100	70	100	60	100	50	100	60	100	35

NOTES: Percentages in columns 00 through 65 are total CU utilizations, that is, the sum of CU utilizations by the file transfer(s) and by other (open-loop) activity in amount specified for column.

Percentages in columns "FT" indicate CU utilization by file transfer(s) only, that is, the total CU utilization minus the open-loop utilization as specified in its column heading.

Figure 13. Control Unit Utilization by File Transfers

In CUT mode, only the 3X74 data character codes (96) can be transferred through the subsystem control unit. For a file with code points outside this set, like

a file with binary data, the file transfer programs in both the host and the workstation must have translate facilities to make sure that only legal character codes are being transmitted through the control unit at all times.

Algorithms to accomplish this can be designed in many different ways. While the length of the file being transferred must always increase when its code point set exceeds the permissible CUT mode set, some algorithms are more successful than others in minimizing this expansion.

An example of a simple but rather poor approach is to divide each file byte into two 4-bit nibbles, concatenate each with four bits to ensure two legal code points for transmission, and, upon receipt, reassemble both into the original file byte. This doubles the number of bytes passing through the control unit, irrespective of whether the file is a text or binary file.

In DFT mode such algorithms are not necessary because any of the 256 possible codes in an 8-bit byte is allowed to pass through the control unit using the WSF command. While this particular aspect would seem to favor DFT over CUT mode, this difference may not play a significant role, depending on file content and the efficiency of the algorithm used by a particular program product.

Note: To perform file transfers on IBM Personal Computer workstations (with an IBM 3278/79 emulation card) attached to a 3174 control unit, and operating in CUT mode, make sure that the File Transfer Aid (customizing question 125) is set to 1.

FILE TRANSFER BENCHMARKS

Five benchmark files with lengths (F) of 0.4, 2, 8, 27, and 50 kbytes, and containing text characters only, were used for obtaining the data in this document

Measurements with different file lengths generally confirm that actual file transfer time (B x F) is proportional with file size within a few percent.

EFFECT OF THE WORKSTATION FILE DEVICE

The workstation contribution "Bw" to file transfer performance depends on the available amount of RAM buffer area, the file device (fixed disk or diskette drive), how fully it is loaded, and the extent of its load fragmentation.

File blocks are always transferred into or out of the main storage (RAM) of the workstation. When the available storage area can contain many of these blocks, the performance effect of rotating file device accesses will be slight, and thus the transfer rate will be near maximum. (This is sometimes called RAM file transfer.) In this connection, note that in some workstations, the 3270-PC for example, it is possible to control this buffer area size by setting a value in the CONFIG.SYS file. (One might try BUFFERS=10 as a starting point.)

When the RAM area can store relatively few blocks, frequent access to the file device is required. This can slow down average file transfer rate, depending on whether file access time overlaps processing in the host and control unit, the file device characteristics, how fully the file is loaded, and how fragmented its remaining storage space is.

For fixed (hard) disks, average file transfer rates may approach RAM-file rates. Transfers to diskette drives will be slower: how much slower depends on many, mostly intractable, factors.

AN EXAMPLE OF HOW TO ESTIMATE FILE TRANSFER PERFORMANCE

This example illustrates how a file transfer time estimate can be obtained with the information provided in this chapter and data on the file transfer characteristics of a 3X74 with a specific workstation. The table in Figure 14 lists the data on a 3194 workstation, for use in this example.

File Transfer Performance of 3194 Display Stations
Estimated Time: A + (B x E) x F seconds

CU	WS	FDV	Size/Mode	Dir	RU/Blk	A	B	FTR	Bc/B
3174 SNA	3194	Dsk	2048/DFT	Rec	1024	4.3	0.22	4.5	0.13
				Snd	256	6.1	0.35	2.9	0.17
3174 nSN			2048/DFT	Rec	2048	2.3	0.32	3.1	0.09
				Snd	2048	2.9	0.39	2.6	0.09
3174 nSN			1920/CUT	Rec		15	0.44	2.3	0.10
				Snd		15	0.45	2.2	0.10

NOTES:

A	Initiation + Termination time, in seconds
B	Seconds per kbyte (= Bh + Bc + Bw)
E	Expansion factor (no dimension)
F	File size, in kbytes
FTR	File Transfer Rate (= 1/B), in kbytes per second
Bc/B	Ratio used for determining E
CU	Subsystem control unit
WS	Workstation
FDV	File Device: Fdk - Fixed disk; Dsk - Diskette
Size	Size of file block being transmitted, in bytes
Mode	CUT mode or DFT mode
Dir.	Direction of file transfer: Rec (Receive), or Snd (Send), or R/S (Rec/Snd)
RU/Blk	Size of RU (SNA) or Block (non-SNA), in bytes

Figure 14. 3194 Version 1 File Transfer Performance Data

Let us assume that files of 80 kbytes (= F) are being transferred to 3194 workstations attached to a 3174, operating in DFT mode in a non-SNA environment, and that the utilization of the control unit by MFI (open-loop) transactions is 20 percent (for example, twenty CUT workstations with a full MFI transaction load).

From the appropriate line in Figure 14 we obtain for "non-SNA DFT Receive":

A = 2.3 seconds
B = 0.321 seconds per kbyte (FTR=3.1 kbytes per second)
Bc/B = 0.094

From Figure 12 on page 61 we copy, for Bc/B=0.100 in the "20" (percentage) column, the E-values for one, two, three, and four concurrent transfers below. Adjusting for E=0.094 by means of linear interpolation between these values and 0.000 for Bc/B=0 we obtain:

Bc/B=	1 XFR	2 XFRs	3 XFRs	4 XFRs
0.000 ---->	E=1.000.	E=1.000	E=1.000	E=1.000
0.100 ---->	E=1.026	E=1.041	E=1.060	E=1.083
0.090 ---->	E=1.023	E=1.037	E=1.054	E=1.075
Tf, seconds	28.6	28.9	29.4	29.9
U total, %	30	37	44	50

Substitution of A, B, E (for a single transfer), and F in the expression for Tf yields:

$$Tf = 2.3 + (0.321 \times 1.023) \times 80 = 28.6 \text{ seconds}$$

The Tf-values for two, three, and four concurrent file transfers are listed above. The associated approximate total utilizations, obtained from Figure 13 on page 62, are listed as well.

EFFECT OF FILE TRANSFERS ON MFI TRANSACTIONS

The 3174 utilizations listed above show that four concurrent file transfers increase this utilization from twenty to about fifty percent.

The information in Figure 7 on page 44 shows that, for example, for the A-1200 benchmark the subsystem response of a CUT-mode 3191 display station is 121 milliseconds at a 20 percent 3174 utilization (non-SNA). At 50 percent utilization, this average response is estimated to increase to 182 milliseconds.

CHAPTER 9. REMOTE SUBSYSTEM CONTROL UNITS

For remote 3X74 control units, the speed and type of the telecommunication link rather than the control unit determines performance.

Because transmission time depends, among others, on data stream length, it is important to minimize data stream **lengths** where possible, to minimize response times, and to maximize a link's capacity. For example, some high speed printers have compression or compaction options, while some 3270 commands and orders can sometimes help reduce data stream length. (A row of dashes across a page, for example, can be compressed to four characters by using the Repeat-to-Address order.)

In both SNA and non-SNA environments, a remote 3X74 control unit starts the processing or pass-through of a data stream as soon as the first RU segment or bytes of a BSC block have been received. Because the processing or pass-through rate almost always exceeds the arrival rate, the subsystem needs relatively little time to finish up once the transmission is complete, usually in the order of 10 to 20 milliseconds. Thus, for a given link speed and protocol, you should not expect significant performance differences between various remote 3X74 control unit models.

To get an idea of the data transfer capacity of a link, first divide the bit rate (in kbps) by 8 to obtain the equivalent of this link rate in kbytes per second. The actual **data** rate that can be achieved is less than that because data streams are divided into RUs or blocks, and are therefore augmented with bytes for flags, headers, and checking, or control characters in the case of BSC. Furthermore, the protocol involves polling, acknowledgments, error, and recovery messages introducing propagation and modem delays.

To account for these factors in an estimate of the capacity of a link, assume that the aggregate utilization of a full duplex SDLC or BSC line by data bytes cannot exceed fifty percent of its speed: thirty percent for a half duplex BSC line. (Note that even on a line with full duplex capability, 3X74 control units support half duplex operation only.) Therefore, a 56 kbps SDLC line would be estimated to sustain an average data transfer rate of 0.50 x (56/8) = 3.5 kbytes per second.

Actual data rates may be higher, and do in fact occur, although line utilization on the basis of data content will never reach 100 percent for the reasons explained above. The utilization percentages in Figure 21 on page 85 illustrate this point (although they are somewhat higher than they would be on a data rate basis because headers, etc. are included). Note, however, that these numbers also illustrate the considerable increases in average response time associated with high utilization.

With enough devices attached to a remote control unit, chances are that, at times, a link will be near 100 percent utilized, causing device operations to heavily interfere with each other, resulting in high, unpredictable response times. For example, the operation of a high speed printer, continuously requiring a substantial share of the line capacity, may be frequently interrupted by the operations of other devices. The same holds true for file transfer operations. MFI operations on the same remote control unit may be

frequently delayed as a result of medium or high speed printing, or file transfers. They will not be inhibited altogether.

Note that all file transfer performance data in this document apply to local 3X74 subsystems. For file transfers to and from remote subsystems, Bc (Chapter 8) will increase very significantly! (Subsystem times for transferring a block of file data may increase from something in the order of 0.1 second for a local subsystem, to seconds, depending on link speed and other details.) Not only will the transfer times for a given amount of data become much longer, but the average number of concurrent transfers will probably increase as well, slowing things more. Caution is advised.

The recommendation for long data streams to remote subsystems is, as for local attachment, to use of the largest feasible RU size with a pacing of two (or more), or comparable block sizes. This is a general rule; for a particular case it may pay to do some fine tuning.

Long data streams to DFT workstations on remote non-SNA control units (BSC) in an NCP environment may perform more poorly than necessary, if the available means for releasing the telecommunications line during data stream processing in the workstation are not used.

BSC CONSIDERATIONS

When attaching 3278/79 displays to a 3174-1R or -51R rather than a 3274-41C or -61C control unit, no noticeable performance differences should be expected for most data streams. Exceptions may be high speed links and very complex data streams, such as associated with graphics. Remote configurations were not specifically modeled for this bulletin.

A 3X74 will accept message sizes up to 7 Kb of data that ends with an ETX, and up to 3.5 Kb for blocked data ending with DLE,ETB. (See Appendix E for an explanation of the abbreviations.) Blocked data are recognized only if the message is sent in transparent mode. If the data is not in transparent mode, a 3X74 processes an ETB as an ETX, meaning end of text. If the data count exceeds these limitations, the 3X74 will send EOT, and set OP Check status.

The 7 Kb maximum message size is only provided as a migration aid. Line errors may cause frequent retransmission of large messages/blocks, thereby degrading performance and overall throughput. Maximum message/block size should not exceed 3.5 Kb, while for a given environment optimum sizes should be obtained from network studies.

For 3290 display stations, or other workstations operating in DFT mode attached to a 3X74 using the BSC TP protocol, some new facilities have been provided for transmitting lengthy outbound data streams to accommodate applications that use a large capacity screen on the 3290, or that use workstations for display of vector or image graphics.

- WACK response to outbound blocks ending in DLE,ETX or ETX
- ACK/WACK responses to outbound text blocks ending in DLE,ETB
- Maximum outbound message size of 7 Kb unblocked data ETX.

A 3X74 control unit passes messages through to or from DFT devices rather than processing (or generating) their content as with 3278/79 displays. By returning a WACK response after an outbound block ending with DLE,ETX or ETX is received and passed on to a workstation, the control unit becomes free to service others. The control unit will accept either EOT or ENQ in response to WACK. If the host sends EOT, ending status DE or 'Error' will be reported asynchronously to either a General or Specific Poll. The ENQ response causes ending status to be handled synchronously, meaning the host will be in a WACK-ENQ loop until the device has completed processing.

To prevent error logging of possible line timeouts, a 3X74 will return WACK in response to DLE,ETB immediately if no buffers are available in the workstation to receive the next outbound transmission. The control unit and device remain selected on the line (line 'busy'). The host sends ENQ in response to WACK. The control unit waits two seconds before returning the second or subsequent WACK. When buffer resources become available, a 3X74 will return ACK.

Without WACK/EOT support in the host, the line and control unit are busy until data stream processing in the workstation is complete. For certain data streams, TP-line utilization may increase by 100 percent or more.

Messages using blocking DLE,ETB should be added to the network only after careful planning. As noted above, the line and control unit are busy until the ETX is received and the data stream is off-loaded to the DFT workstation. For example, if a 21 kbyte 'blocked' message is sent, a 9600 bps line and the control unit would be busy for approximately 25 seconds (43 seconds for 4800 bps). This is unacceptable to other terminals on the line. If two terminals requested large screens, the time could approach 50 seconds.

It is recommended to use the WACK support in the NCP, where available. For a remote BSC 3X74 with 3290s attached to a line controlled by an NCP without WACK support, line performance will be less efficient when sending to 3290s than when sending to 3278/79s. A 3X74/3278 responds to an outbound data stream with an ACK, which causes the NCP to move on to its next activity. A 3X74/3290 responds to an outbound data stream with a WACK, which causes the NCP to issue ENQ back to the 3X74. The line and control unit will remain selected until 3290 processing completes and the 3X74 returns an ACK. Depending on message content, message length, and line speed, line throughput degradation could be severe.

CHAPTER 10. IBM 3174 SUBSYSTEM WITH IBM TOKEN-RING NETWORK

This chapter contains the entire contents of the seventh edition of this document, ZZ20-4167-6, which was electronically distributed on HONE exclusively.

TOKEN-RING NETWORK CONFIGURATION

The 3174 local Model 1L, and remote Models 1R, 2R, 51R, and 52R Subsystem Control Units, when equipped with the IBM Token-Ring Network 3270 Gateway (#3025), and customized for SNA, can attach an IBM Token-Ring Network (TRN). This token-ring can attach several 3174 Model 3R and 53R control units, as well as other workstations. See Figure 15. (The use of "-3R" implies "-53R" as well, unless expressly excluded.)

```
                                                         |
HOST CHANNEL ====||========      TELECOMMUNICATIONS |
             ┌───┴┴───┐               LINK      ┌───┴────┐
             │  3270  │                         │  3270  │
             │Gateway │                         │Gateway │
             │ (SNA)  │          or             │ (SNA)  │
             │3174-1L*│                         │  3174  │ -1R*, -2R*,
             └┬┬┬───┬┬┘                         └┬┬┬───┬┬┘ -51R*, or -52R*
              |||   ||                           |||   ||
     3270 Devices   ||                  3270 Devices   ||
┌───────────────────┘└──────────────────────────────────────────────────┐
│                                                                        │
│                       Token-Ring Network                               │
│                                                                        │
└────┐┌─────────┐┌────------────┐┌───────------──┐┌─┐┌─┐┌──--─┐┌─┘
  ┌──┴┴───┐  ┌──┴┴───┐       ┌──┴┴───┐            ||  ||  ||    ||
  │Ring-  │  │Ring-  │       │Ring-  │         Other workstations
  │attached│ │attached│ UP TO│attached│
  │ 3174-3R│ │ 3174-3R│ 12 CUs│ 3174-3R│
  └┬┬┬───┬┘  └┬┬┬───┬┘       └┬┬┬───┬┘
   |||   |    |||   |         |||   |
3270 Workstations and printers ───>
```

Figure 15. IBM Token-Ring Network Configurations

A 3174 Subsystem Control Unit with the 3270 gateway feature installed is referred to as a **gateway control unit** and is indicated with an asterisk (*) to differentiate it from 3174s without this feature: 3174-1L*, or 3174-1R*, 2R*, 51R*, and 52R*. For example, 3174-1L*/TRN/3174-3R/3278 specifies a Token-Ring Network attached to a local 3174-1L with 3270 gateway, and 3174-3R control units with 3278 display stations attached to the ring. A 3278 directly attached to a local 3174 subsystem control unit with a gateway feature would be shown as "3174-1L*/3278".

Configuration Support-S Release 2 not only makes the 3270 gateway perform somewhat better than with Release 1, but supports, in addition, the use of the 3270 gateway in the 3174 Models 1R, 2R, 51R, and 52R subsystem control units for connection to an SDLC telecommunication link with a speed of up to 64 kbps.

This edition includes the following new information with respect to the Release 1 data in Chapter 5 of the sixth edition of this document, ZZ20-4167-5, and in **Performance of IBM 3174 Subsystem Control Units,** ZZ27-2671-0.

- Improved 3270 gateway performance data as a result of using Configuration Support-S Release 2.
- Performance of IBM Personal Computer workstations directly connected to a 3174-attached token-ring.
- Token-Ring Network performance for 3174 Models 1R, 2R, 51R, and 52R subsystem control units with a 3270 gateway, when connected to 56, 19.2, and 9.6 kbps telecommunication links.

The 3270 gateway in local and remote 3174 control units function as a 'pass-through' multiplexer for Token-Ring Network traffic to and from a host through a local channel or a telecommunications line. Because the 3270 gateway does not care about source or destination device type, the traffic rates and utilizations apply to a wide range of attachments. Be aware that some 3174 customization parameters, such as the maximum frame size allowed on the ring, may affect 3270 gateway/token-ring performance.

One of the following criteria determines the maximum number of ring-attached control units and workstations that can have access to the host:

- The maximum number of SNA Type 2.0 physical units defined to VTAM as being capable of host access
- Performance of the 3270 gateway when using a local 3174 Model 1L* control unit, or
- Performance of the telecommunications line when using a remote 3174 Model 1R*, 2R*, 51R*, or 52R* control unit.

The maximum number of Type 2.0 physical units (PUs) supported on a ring depends on the amount of storage installed in the gateway control unit. (The maximum is 140 sessions with two #1012 storage expansion features: see the sales pages.) A 3174 control unit or workstation attached to a ring represents one physical unit. In turn, a 3174 Model 3R control unit is capable of supporting up to 128 logical units, while Model 53R supports up to 76 LUs.

From a performance point of view, the number of 3174 control units that you may want to attach to a ring is likely to be more limited. Transactions between a workstation on a 3174 ring-attached control unit and the host pass through four subsystem layers, as shown in Figure 15 on page 71.

- The host channel, or the telecommunications link
- A local or remote 3174 subsystem control unit with the 3270 gateway function for attachment to a Token-Ring Network

- The Token-Ring Network
- The ring-attached 3174 Model 3R and 53R control units, capable of attaching up to thirty-two and sixteen 3270 devices, respectively.

When an operator at a workstation connected to a ring-attached control unit initiates an MFI transaction, for example, an inbound message first passes through this control unit, then through the Token-Ring Network and the gateway control unit, on to the host. When the host delivers its response, an outbound message passes through these subsystem layers in reverse order. In addition, there are acknowledgment messages.

For device/host communication, 3270 gateway traffic equals ring traffic. Because for a given amount of traffic in a 3270 gateway, utilization of the ring is about one fourth of the gateway utilization, some peer-to-peer traffic over the ring is not likely to have a significant effect on the performance data in this chapter.

When a local 3174 control unit is attached via an IBM 3044 Fiber-Optic Channel Extender Link, be aware that the maximum amount of traffic through the 3270 gateway may be constrained by this link rather than by the gateway itself, especially for longer link lengths. No performance data is available for 3174 subsystems using this link. (See **IBM 3044 Fiber-Optic Channel Extender Link,** GA22-7097, for more information.)

ASSUMPTIONS

For the performance assessment of 3270 gateways in **local 3174s,** a 3174-1L*/TRN/3174-3R/3278-2 subsystem with twelve 3174 Model 3R control units (or up to twenty-four 3174-53R units) on the token-ring is used. The ring-attached control units are capable of connecting up to 384 workstations and printers. (Raising the number from ten to twelve is based on using 3174 Configuration Support S Release 2 rather than Release 1.) The combined total message traffic generated by all workstations is limited by the 3270 gateway.

To assess performance of a gateway in a **remote 3174,** a 3174-1R*/TRN/3174-3R/3278-2 subsystem is used. The number of workstations (ring-attached control units) that can be supported is considerably lower, because the gateway traffic is now limited by the speed of the telecommunications link.

An average workstation transaction load of 6.7 type A-1200 MFI transactions per minute was adopted (400 transactions per hour, or an average interarrival time (IAT) of nine seconds, independent of host or subsystem response times). This amounts to 214 transactions per minute (=CTR) for 32 active workstations on a control unit, and 2568 transactions per minute (=RTR, ring transaction rate) for twelve 3174 control units on a token-ring.

The 3270 gateway in a Model 1L* limits the combined traffic of all ring-attached control units and workstations with the host. Therefore, when ring-attached control units contribute less than the assumed number (214) of transactions per minute, because fewer of their workstations are active and/or have lower average transaction rates, more control units/workstations can be attached to the token-ring, as long as the 3270 gateway capacity is not exceeded. On the other hand, the number of usable control unit ports may be fewer when many have high data rate requirements, such as for frequent file transfers, graphics, image, and/or high-speed printing.

The performance of a file transfer operation does not only depend on 3X74 subsystem characteristics, but also on the processing times of the support code in the workstation and the host (and network, if any). Because file transfer is a closed-loop operation, fast turnaround times in the workstation and the host favor high data transfer rates, but raise subsystem utilization during the duration of the transfer. Conversely, delays slow transfers, but decrease subsystem utilization.

The file transfer performance data in this document is based on the IBM 3270 Personal Computer AT workstation in DFT mode with its Release 2.1 control program. Assumed average turnaround times in workstation and host are 85 and 50 milliseconds per 2 Kb block of data, respectively.

Selection of the **maximum frame size of 2042 bytes** on the token-ring is recommended, and was used for obtaining the performance data. (See **IBM 3174 Subsystem Control Unit Customizing Guide,** GA23-0214: respond to configuration questions 941F (1L* and R*) and 380 (3R and 53R) with "2042" and "3", respectively, (for outbound frames), and to question 382 for Models 3R and 53R with "2042" (for inbound frames)). Pertinent configuration questions have been summarized in Figure 16 on page 75.

Small token-ring frame sizes for large messages increase the 3270 gateway utilization substantially, because frame overhead processing increases due to the increased number of frames required to transmit a given size (large) message. As a result, traffic-handling capacity of the gateway may be reduced by as much as fifty percent, depending on the application.

The "Maximum-In" and "Maximum-Out" (transmit window size) parameters in Figure 16 specify the number of transmits before waiting to receive a token-ring level acknowledgment. The performance data was obtained with a setting of "2" for outbound messages (questions 381 and 941W), and "1" for inbound messages (question 383). For a single token-ring, question 941W = 2 speeds up transmission because waiting for acknowledgments is reduced, and keeps buffer resources in the 3270 gateway from being depleted. Selecting "1" for question 383 ensures that each received frame is acknowledged (little is added to token-ring congestion by the acknowledgment).

Question	Parameter Selection	Designation
	Outbound:	
941F:	0 1 2 3	Transmit I-frame size
=380:	0265 0521 1033 2042	Maximum receive I-frame size
941W:	1-7 1-7 1-4 1-2	Maximum out (transmit window size)
381:	equal or less than 941W	Maximum in
	Inbound:	
382:	select from 0265 to 2042,	Maximum transmission I-frame size
	when 0265-1033, 1034-2042	
383:	1-7 1-4	Maximum out

NOTES:

- Questions 941F and 941W pertain to the 3174 Model 1L* and R*, while 380, 381, 382, and 383 apply to Models 3R and 53R.
- All frame sizes in bytes, includes TH and RH headers
- Set same frame length for Questions 380 and 941F
- Do not set 941W to be less than 381
- Performance data in this document is based on "3", "2042", and "2042" for questions 941F, 380, and 382, respectively; "2" for questions 941W and 381; and "1" for question 383.
- Defaults: Q380="2042", Q381="1", Q382="521", QQ383="2", Q941F="3" (2042), and Q941W="2".

Figure 16. Token-Ring Network Parameter Selection

At the SNA level, one SNA response (+DR) has been assumed to follow the 'write' operations in both type A-1200 benchmarks and file transfers. While an SNA response message does not add much to response time, it does significantly

increase 3270 gateway utilization, thereby decreasing the maximum amount of traffic that the 3270 gateway can handle.

Performance data is presented in tabular form rather than as curves on charts. This data can, of course, be used to construct a graph. Intermediate values may be approximated by linear interpolation.

Response times, in milliseconds, include the response of all elements in the subsystem, that is, workstation, a ring-attached control unit, the token-ring, and the 3270 gateway, but **not** the host turnaround time. Utilizations are in percent, unless indicated otherwise. The data does not include contingencies.

TOKEN-RING NETWORK ON LOCAL 3174* CONTROL UNITS

The data in Figure 17 compares, for various A-1200 transaction rates, the response times of a local 3174-1L*/TRN/3174-3R/3278 subsystem with twelve Model 3Rs, with the responses of 3278s directly attached to a 3174 Model 1L* and a 3274 Model 41A.

Note that 6.7 A-1200 transactions per minute on a 3278-2 display station cause an average utilization of 0.17 and 0.95 percent in the 3270 gateway and the 3174-3R control unit, respectively. For a 3174-3R with 32 workstations (214 transactions per minute), these utilizations would be 5.53 and 30.4 percent, respectively.

		Local Subsystem Response Times as a function of A-1200 transaction rates							
1a	A-1200 trans/minute in -3R	6.7	40	80	120	160	200	240	280
1b	WSs doing 6.7 trans/minute	1	6	12	18	24	30	–	–
2a	3174-3R utilization	.95	5.7	11.4	17.1	22.8	28.5	34.2	39.9
2b	3270 gateway utilization	.17	1.04	2.08	3.12	4.15	5.18	6.25	7.15
3a	3174-3R/1 (RTR=0)	117	122	129	137	146	155	165	177
3b	3174-3R/6 (RTR=1070)	121	126	133	141	150	159	169	181
3c	3174-3R/12 (RTR=2360)	137	142	149	157	166	175	185	197
4	3174-1L* (RTR=0)	109	114	121	129	139	149	162	175
5	3274-41A	192	213	245	293	369	500	-	-
6a	3174-3R/6, with 6 FTs	169	174	181	189	198	207	217	229
6b	3174-3R/6, with 12 FTs	227	232	239	247	256	265	275	287

NOTES:

- Response data is in milliseconds, and control unit utilizations are in percent.
- Data for 3174-1L* with token-ring with twelve 3174-3R control units attached, each with 32 3278-2 display stations; no traffic from workstations directly attached to 3174-1L*
- 3174-3R/n: n= number of 3174-3R ring-attached control units
- RTR Ring transaction rate, in transactions per minute
  FT File transfer
  WS Workstation

Figure 17. 3174-1L*/TRN/3174-3R/3278 Subsystem Response Times

Entries in the table are explained below:

- Line 1a: Indicates transaction load in a ring-attached control unit, in type A-1200 transactions per minute.

- Line 1b: Specifies transaction load in Line 1a in terms of active workstations, each transacting 6.7 A-1200 benchmarks per minute.
- Line 2a: Lists the utilization of the ring-attached control unit as result of the transaction load on Line 1a, in percent.
- Line 2b: Lists the associated utilizations of the 3270 gateway.
- Line 3a: Response time on a single ring-attached control unit as a function of its load, with no additional ring traffic (designated "3174-3R/1").
- Line 3b: Same as Line 3a but with ring traffic of five other ring-attached control units, that is, 1070 (= 5x32x6.7) additional transactions per minute (designated "3174-3R/6"). (The difference with Line 3a represents the added gateway/ring delay as a result of the additional load.)
- Line 3c: Same as Line 3b but with traffic from eleven other control units with 2360 additional transactions (designated "3174-3R/6").

  These response times, and those in Line 3b, exhibit constant differentials with the response times in Line 3a, due to the additional traffic in the 3270 gateway/token-ring. (Subsystem response estimates for intermediate gateway/ring traffic conditions can be obtained by interpolation.)
- Line 4: Response of 3278s directly attached to a 3174 Model 1L* equipped with a 3270 gateway, but no traffic through the ring/gateway. (This response time is about 3 milliseconds more than for a Model 1L without the 3270 gateway option). Even so, subsystem response time is less than for 3274 control units, and exhibits a slower rise with increasing transaction rate.
- Line 5: Response times of a 3278 attached to a 3274 Model 41A.
- Line 6a: Same as Line 3b, but with six file transfers added to the gateway/ring traffic (with a combined data rate of about 39 kbytes per second, or about 6.5 kbytes per second for each transfer).
- Line 6b: Same as Line 3b, but with twelve file transfers added (with a data rate of about 58 kbytes per second, or about 4.8 kbytes per second per transfer).

PC WORKSTATIONS DIRECTLY ATTACHED TO TOKEN-RING NETWORKS

IBM Personal Computer XT and AT workstations, equipped with the appropriate adapter and PC 3270 Emulation Version 3 support, can directly attach to a Token-Ring Network. Such a workstation is defined to VTAM as one physical unit (PU).

In an MFI host application, such a workstation may be used for emulating a 3270 DFT mode terminal. Lines 3a through 3c in Figure 18 provide its subsystem response times at three levels of ring traffic: none, intermediate, and high.

		Local Subsystem Response Times as a Function of Number of Workstations							
1a	A-1200 transactions/minute	6.7	133	267	400	533	667	800	933
1b	WSs doing 6.7 trans/minute	1	20	40	60	80	100	120	140
2	3270 gateway utilization	.31	6.1	12.2	18.3	24.4	30.5	36.6	42.8
3a	Subsystem resp with RTR=0	397	398	399	401	405	409	412	415
3b	With RTR=1070 A-1200/min	398	401	412	424	432	441	452	463
3c	With RTR=2360 A-1200/min	458	465	487	552	743	1304	-	-

NOTES:

- Response data is in milliseconds, and control unit utilizations are in percent.
- 3270 Personal Computer AT workstation with PC 3270 Emulator Version 3, directly attached to Token-Ring Network; no traffic from workstations directly attached to 3174-1L*
- RTR Ring A-1200 transaction rate generated by 3174-3R control units, in transactions per minute
- WS Workstation

Figure 18. 3174-1L*/TRN/3270PC-AT with PC 3270 Emulator Version 3

Entries in Figure 18 are explained below:

- Line 1a indicates total A-1200 transaction load generated by all ring-attached workstations.
- Line 1b specifies transaction load in Line 1a in terms of active workstations, each transacting 6.7 A-1200 benchmarks per minute.
- Line 2: Lists the utilization of the 3270 gateway resulting from the transaction load of Line 1a, in percent.
- Line 3a: Response time at a single ring-attached workstation, with no additional ring traffic, in milliseconds.

- Line 3b: Same as Line 3a, but with RTR=1070, that is, with 1070 A-1200 transactions through the 3270 gateway/token-ring, originating at 3174-3R control units.

- Line 3c: Same as Line 3b, but with RTR=2360.

Note: For 3270 Personal Computer AT workstations directly attached to the token-ring, the 3270 gateway utilization is higher because the maximum size of transmitted frames is 256 bytes rather than 2042 bytes. This increases the overhead associated with the transmission of a long data stream over the ring.

3270 GATEWAY CAPACITY

The data in Figure 19 focuses on the traffic handling capability of the 3270 gateway in a 3174 Model 1L* control unit.

Because the service times of messages in the 3270 gateway/token-ring combination are an order of magnitude less than their service times in a ring control unit, average gateway delays associated with a given utilization (of the gateway control unit) are also much smaller. Gateway utilizations over 65 percent have, therefore, a smaller effect on subsystem response times than similar utilizations in a control unit would have. For this reason, the 3270 gateway capacity data in Figure 19 has been extended to 85 percent gateway utilization.

		3270 Gateway Capacity							
1	3270 gateway utilization	1	15	30	45	55	65	75	85
2	Gateway–ring delay, ms	18	22	24	28	32	42	56	115
3	RTR traffic, A–1200/min	0	581	1161	1743	2130	2517	2904	3292
4	FT traffic, in kbytes/s	0	18	39	55	70	84	95	105
5a	3174-3R CTR=6.7 RT, ms	117	119	121	123	129	137	149	209
5b	3174-3R CTR=214 RT, ms	158	160	162	164	170	178	190	250

Figure 19. 3174/TRN/3174/3278 Response Times / 3270 Gateway Capacity

Note line-by-line explanations of the entries in the table below:

- Line 1: Lists a range of 3270 gateway utilizations, in percent.
- Line 2: Indicates the contribution of 3270 gateway plus token-ring network delay to A-1200 subsystem response time (inbound plus outbound).
- Line 3: Specifies the ring transaction rates, in A-1200 transactions per minute, that will bring about the 3270 gateway utilizations in Line 1.
- Line 4: Indicates the combined file transfer rates, in kbytes per second, associated with the 3270 gateway utilizations in Line 1.
- Line 5a: Subsystem response obtained at a ring-attached 3174-3R/3278 control unit with CTR=6.7 transaction load.
- Line 5b: Subsystem response obtained at a ring-attached 3174-3R/3278 control unit with CTR=214 transaction load.

Because ring utilization is so much less than 3270 gateway utilization, some peer-to-peer traffic over the ring is not expected to substantially affect this information.

FILE TRANSFER

As noted earlier, file transfer performance data is based on the IBM 3270 Personal Computer AT workstation attached to a Model 3R control unit, and with its Release 2.1 control program customized for DFT mode. The assumed average turnaround times, including acknowledgment messages, associated with the outbound transfer of a 2 Kb block of data were 85 milliseconds for workstation, and 50 milliseconds for the host. Therefore, Bw=0.043 and Bh=0.025 seconds/kbyte, respectively.

Maximum file transfer rates are achieved with minimum gateway/ control unit response time, that is, low gateway and control unit utilizations. When these utilizations increase, file transfer rates will decrease.

Measurements yield a minimum subsystem time, that is, with no ring and control unit traffic, for a 2 Kb block of data of about 115 milliseconds (Bc=0.057 seconds). Using the expression:

$$\frac{1}{B} = \frac{1}{Bh + Bc + Bw} = \text{File transfer rate in kbytes/second}$$

yields, therefore, a maximum outbound file transfer rate estimate for workstations on ring-attached control units:

$$\frac{1}{0.043 + 0.057 + 0.025} = 8 \text{ kbytes/second}$$

Assuming that Bc increases in proportion with RT, Bc can be estimated for various load conditions of the gateway and the ring control unit on the basis of the response times in lines 5a and 5b in Figure 19 on page 81. For example, for a 65 percent gateway utilization and a loaded Model 3R control unit, Bc can be assumed to be about 52 percent (178 versus 117) more, that is, Bc=0.087 seconds, yielding a file transfer rate estimate of about 6.4 kbytes per second.

Corresponding maximum inbound rates are expected to be somewhat less.

Based on utilization numbers obtained from modeling file transfer in the subsystem, it is recommended that an 105 kbyte/second aggregated rate of concurrent file transfers through the 3270 gateway not be exceeded. This corresponds with a gateway utilization of about 85 percent, and a ring utilization of about 18 to 20 percent.

EFFECT OF 3174-1L* ACTIVITY ON THE GATEWAY CONTROL UNIT

The data in Figure 20 lists utilizations of the 3174 Model 1L* control unit as a result of certain transactions when performed on a workstation attached to a Model 3R control unit, as opposed to being performed on a workstation directly attached to the gateway control unit.

These numbers show that a given transaction at a directly attached workstation brings about a much larger Model 1L* utilization than when it is performed at a workstation transacting through the token-ring.

Transaction description:	Utilization of 3174–1L*, in percent: Through 3174–3R & gateway	Directly attached
6.7 A–1200 transactions/second, in CUT mode	0.17	0.95
Associated keystroke processing, average	0	0.50
6.7 A–1200 transactions/second, in DFT mode	0.17	.60
Associated keystroke processing	0	0
File transfer at 1 kbyte/second	0.81	3.5
Typematic NL key actuation (on CUT display)	0	17

Figure 20. 3174-1L* Utilization by Various Transaction Types

With respect to directly attached CUT mode workstations, note the very significant utilization as a result of processing New-Line keystrokes at a 10/second typematic rate. (NL keystroking is used as an example because it requires complex processing, with precedence over much processing in support of the Token-Ring Network.) With several of these going on at the same time, the effect on token-ring performance may be quite noticeable.

With workstations operating in DFT mode, control unit utilization is **not** affected, because keystrokes are processed in the workstation.

Because the utilization of the gateway control unit is critical for satisfactory token-ring network performance, one should be selective in the attachment and use of workstations to Model 1L* control units. It is recommended to plan for a total gateway control unit utilization of not more than 85 percent. This caveat is less important for remote 3174-R*s because the SDLC telecommunications link limits 3270 gateway throughput to well below ten percent of its capacity.

The delay associated with different utilizations can be obtained from Figure 19 on page 81.

TOKEN-RING NETWORK ON REMOTE 3174* CONTROL UNITS

A Token-Ring Network can be attached to a 3174-1R*, 3174-2R*, 3174-51R*, or 3174-52R* Subsystem Control Unit, which, in turn, is connected to a host with a SNA/SDLC telecommunications line. (Installation of the IBM Token-Ring Network 3270 Gateway optional feature (#3025) and Configuration Support-S Release 2 is required.) The SDLC link may operate in data half-duplex mode on duplex or half-duplex nonswitched facilities at speeds up to 64 kbps.

The performance of transactions with the host over these token-rings is dominated by the speed of the telecommunications link, and is affected by several other of its characteristics as well. Even for the highest line speed and the most favorable conditions, the 3270 gateway utilization will remain less than ten percent. The 3174-R*/TRN/3174-3R subsystem itself will, therefore, exhibit close to single-thread response, that is, somewhere in the range of 130 to 170 milliseconds for a type A-1200 transaction.

The performance data in Figure 21 on page 85 assumes a nonswitched point-to-point duplex 300-mile line of indicated speed, being used in data half-duplex mode (virtually no modem turnaround time), and with no pause in-between poll cycles of PUs on token-ring (NCP allows PAUSE to be set from 0 to 25.5 seconds with 0.1 second increments). Furthermore, the selected RU sizes were large enough to contain the entire inbound or outbound data streams associated with the A-1200 benchmark. The response times include link and NCP delays, but not the turnaround time in the host.

The data is provided for the NCP PAUSE operand set at 0 and 0.1 second. This operand, specified by the user, determines the minimum delay between the start times of poll cycles polling PUs on the token-ring. (NCP allows pause to be set from 0 to 25.5 seconds with 0.1 second increments; the default is 0 seconds). Long polling cycles as a result of many PUs being polled may take longer than a specified pause.

The data shows that the effect of increasing the pause from 0 to 0.1 seconds is slight. With long polling cycles and/or settings of PAUSE, one should expect these response times to increase.

The effect of delays on the efficiency with which a line can be used is commensurate with the line speed. For example, the effect of a given delay (as from propagation, or modem turnaround) is larger at high than low link speeds because the lost 'transmission opportunity' in terms of bits is larger. In other words, while the transmission time of a data stream of a given length becomes smaller with higher line speeds, the percentage of time that the line is not used due to fixed delays increases.

More specifically, the data in Figure 21 on page 85 is affected as follows:

1. By distance -- the length of the line determines propagation delay (use 'miles/15' for links less than 200 miles, and '12 + (miles/150)' for longer lines, in milliseconds (=14 ms for 300 miles)).

2. Whether the line is duplex or half-duplex, and if duplex, whether the operation is in data half-duplex or data duplex mode. Only data half-duplex operation is supported, but use of a duplex line avoids modem turnaround delay because a return transmission can begin immediately. On a half-duplex

line, modem turnaround delays decrease the effective data rate, especially on higher speed lines.

3. With multi-drop lines, polling can introduce delays, depending on how many stations have been polled before the ready-to-transmit station is reached. Thus the number of physical units being polled has an effect.

4. PAUSE delay, although the effect of increasing PAUSE from 0 to 0.1 seconds is slight, one should expect response times to increase with higher settings of PAUSE, and/or many PUs (longer polling cycles).

		Remote Subsystem Response Times as a function of A-1200 transaction rates					
1a	A-1200 trans/minute in -3R	6.7	40	80	120	160	200
1b	WSs doing 6.7 trans/minute	1	6	12	18	24	30
1c	WSs doing 4 trans/minute		10	20	30	40	50
	56 kbps line, with polling						
2a	pause [4] 0: Response	395	426	481	563	701	1066
	Line utilization[4]	16	26	39	52	64	77
2b	pause [4] 100: Response	425	454	506	586	723	1070
	Line utilization[4]	8	19	34	48	61	75
	19.2 kbps line, polling						
3a	pause [4] 0: Response	833	1122	2982	–	–	–
	Line utilization[4]	36	61	90	–	–	–
3b	pause [4] 100: Response	865	1164	2990	–	–	–
	Line utilization[4]	21	53	88	–	–	–
	9.6 kbps line, polling						
4a	pause [4] 0: Response	1582	2072[3]	–	–	–	–
	Line utilization[4]	52	71[3]	–	–	–	–
4b	pause [4] 100: Response	1591	2080[3]	–	–	–	–
	Line utilization[4]	41	61[3]	–	–	–	–

NOTES: [1] Response data and polling pause in milliseconds; control unit utilization in percent.

[2] Data half-duplex operation on full-duplex line of 300 miles

[3] For 20 A-1200 trans/minute in -3R, or 3 WSs doing 6.7 trans/minute

[4] Line utilization includes data, headers, flags, etc.

Figure 21. Performance of SDLC-Link/3174-R*/TRN/3174-3R/3278 Subsystem

The line-by-line explanations for Figure 21 on page 85 are:

- Line 1a: Indicates total A-1200 transaction load of all ring-attached workstations.
- Lines 1b and 1c: Specify transaction load in Line 1a in terms of active workstations, each transacting 6.7 or 4 A-1200 benchmarks per minute, respectively.
- Lines 2a: Response time at a ring-attached workstation with 56 kbps SDLC link, in milliseconds, and link utilization, in percent, with 0-second NCP polling pause.
- Lines 2b: Same as Lines 2a, but with polling pause set at 0.1 second.
- Lines 3a and 3b: Same as Lines 2a and 2b, but for 19.2 kbps SDLC link.
- Lines 4a and 4b: Same as Lines 2a and 2b, but for 9.6 kbps SDLC link.

3270 GATEWAY PERFORMANCE SUMMARY

The principal characteristics of 3270 gateway performance can be summarized as follows:

- For MFI applications, response times at workstations on ring-attached 3174 control units approach the response of workstations attached to local 3174-1L control units (with or without gateway feature).

- Although the traffic handling capacity of a 3270 gateway in a local 3174 exceeds that of a single channel-attached 3174 control unit (roughly, by a factor of two), the fact that many control units usually share a single gateway imposes a lower limit on host traffic per ring-attached control unit.

- It is recommended that you plan for the equivalent of not more than 3300 type A-1200 transactions per minute through the 3270 gateway, resulting in a 85 percent utilization of the 3174-1L* control unit.

- The 3270 gateway capacity, at 85 percent utilization, in terms of file transfer is an aggregate load of about 105 kbytes/second of concurrent file transfers.

- Token-Ring Network utilization associated with 85 percent 3270 gateway utilization is about 25 percent, that is, well below the maximum for obtaining good performance.

- When the aggregate transaction demand from devices attached to 3174 subsystems with a token-ring starts to exceed 3270 gateway capacity, operations will continue without loss of data but response times will increase, and data transfer rates will drop.

- The throughput of 3270 gateways in remote 3174 control units is limited by the traffic handling capacity of the SDLC link.

- Use the maximum permitted frame size for data transmission over the token-ring to a 3174 Model 3R, that is, 2042 bytes. Smaller frame sizes can lower the 3270 gateway traffic capacity. The extent of such a reduction can be substantial in certain cases, depending on the frame size maximum used, message traffic share, and the application.

- In local 3174 gateway control units, the traffic handling capacity of the gateway is lowered by activity generated by directly attached workstations and printers.

  Workstations operating in CUT mode require the control unit to process keystrokes and data streams. (Keystrokes from CUT-mode workstations are processed at a higher priority level than some processing associated with the 3270 gateway.) For this reason, it is preferable to attach DFT workstations with light to medium MFI workloads, rather than CUT workstations (unless sparsely used).

  Avoid the attachment of medium and high-speed printers, workstations with medium to heavy file transfer requirements, and a need for long graphics

and image data streams. **Attach such devices to ring-attached control units instead.**

CHAPTER 11. 3X74-ATTACHED PRINTERS

A variety of printers can be attached to 3X74 Subsystem Control Units. From a performance viewpoint it is of interest:

- Whether the achievable data rate from host through control unit allows the printer to maintain its maximum print speed
- To understand how printer performance is affected by the operations of other devices attached to the control unit
- To know the effect of operating one or more printers on the performance of other devices attached to the control unit.

For occasionally used printers, or those printing not more than a few hundred characters per second, these effects are small and can be ignored. However, for medium and high speed printers in continuous operation, they need to be assessed. The following information is required:

- Printer type, model, print band configuration, maximum print speed, and data buffer size in printer attachment adapter
- Document characteristics, for example, average width of non-blank lines, character content of a line (including underscore), lines skipped
- Subsystem control unit type and model: for example, local or remote, and telecommunication line speed for remote
- The protocol/data stream type: LU type 1 (with SCS or IPDS), LU type 3, or 3270 data stream compatibility (DSC)
- For DSC (and LU type 3), whether an 'Early Print Complete' function is operative.

Printer throughput, the protocol used, and the control unit characteristics determine control unit utilization. For a given data rate and protocol, printer type affects control unit utilization only slightly. Therefore, with no data available on the IBM printer of interest, you may use the data on a comparable printer to estimate the control unit utilization associated with the operation of that printer.

PRINTER CHARACTERISTICS

Printer type and model are principal determinants of maximum print speed. Additional factors affecting speed may be the character set size on the band, belt, or chain for line printers, or the selected quality of the font in dot-matrix printers.

Performance data sheets on IBM printers in the back of this document describe how to estimate the required average data rate to a particular printer on the

basis of printer and document characteristics. For some printers, information on printing speed and throughput may be found in their component descriptions.

DOCUMENT CHARACTERISTICS AND PRINTER DATA RATES

The document characteristics, that is, the printed content of a page, and printer characteristics, interact to yield actual print rate. Whether this rate can be maintained depends on whether the demand for data by the printer is satisfied at all times through the chain formed by the host, data channel or telecommunication link, and control unit. (See Figure 22.) When a printer runs out of data, it halts temporarily. The attendant loss of throughput cannot be recovered.

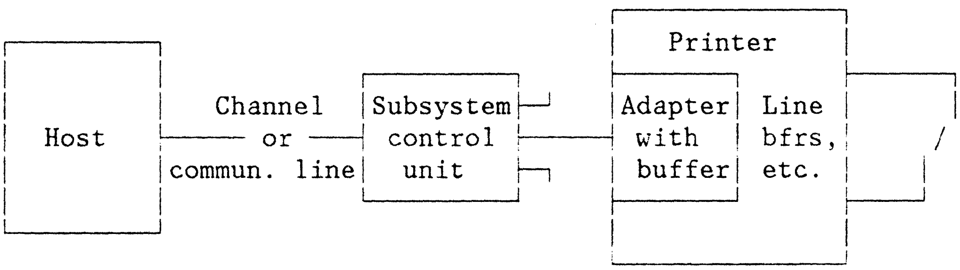

Figure 22. Printer Attachment

For most printers using a print band, the demand for data (in bytes per second) is approximately proportional with the product of print speed (LPM, in non-blank lines per minute) and average print line width on the document, as shown in Figure 23 on page 91. (Some line printers slow down when printing wider lines, even when the demand for data is satisfied at all times.)

Blank lines, underscoring, unusual character patterns, etc. will usually decrease these average data rates depending on printer characteristics, because such operations take extra time.

Since the instantaneous data rate required to keep a printer mechanism operating varies from one moment to the next, and the printer receives data in chunks, the amount of buffering provided in the printer and the data transfer protocol are important. Large buffers will diminish the probability of print interruptions during peaks in demand and/or interruptions in the supply.

SUBSYSTEM CONTROL UNIT TYPE AND MODEL

As far as printer operation is concerned, a major distinction to be made between subsystem control units is whether they are local or remote.

For local control units, it is the performance of the control unit together with the host that determines the upper bounds of the supply of data to a printer.

LPM	Average 36	Line 48	Width 60	in 72	characters 84	96	108	120	132
2000	1.20	1.60	2.00	2.40	2.80	3.20	3.60	4.00	4.40
1800	1.08	1.44	1.80	2.16	2.52	2.88	3.24	3.60	3.96
1600	0.96	1.28	1.60	1.92	2.24	2.56	2.88	3.20	3.52
1400	0.84	1.12	1.40	1.68	1.96	2.24	2.52	2.80	3.08
1200	0.72	0.96	1.20	1.44	1.68	1.92	2.16	2.40	2.64
1000	0.60	0.80	1.00	1.20	1.40	1.60	1.80	2.00	2.20
900	0.54	0.72	0.90	1.08	1.26	1.44	1.62	1.80	1.98
800	0.48	0.64	0.80	0.96	1.12	1.28	1.44	1.60	1.76
700	0.42	0.56	0.70	0.84	0.98	1.12	1.26	1.40	1.54
600	0.36	0.48	0.60	0.72	0.84	0.96	1.08	1.20	1.32
500	0.30	0.40	0.50	0.60	0.70	0.80	0.90	1.00	1.10
400	0.24	0.32	0.40	0.48	0.56	0.64	0.72	0.80	0.88
300	0.18	0.24	0.30	0.36	0.42	0.48	0.54	0.60	0.66
200	0.12	0.16	0.20	0.24	0.28	0.32	0.36	0.40	0.44

Notes: LPM Lines per minute
Data rates in kbytes per second

Figure 23. Data Rate Dependence on Printer Speed and Average Line Width

The maximum data rate capability of the 3174 Model 1L is roughly double that of the 3274 Model 41.

For remote subsystems, the line is the limiting factor, as determined by its speed, block or RU size, pacing, etc. Transmission time almost entirely overlaps control unit operations because a 3X74 starts processing the data as soon as it comes in, keeps pace with the demand, and finishes about 10 to 20 milliseconds after transmission is complete.

PRINTER DATA STREAM AND PROTOCOL

The data stream, that is, the character string (with control and data characters) specifying what and where information is printed, is enveloped by a protocol for transferring it from the host to a printer.

Printer data stream types are:

- SCS -- the SNA character string, with variations depending on type of functions available in the printer.
- 3270 data stream -- uses 3270 data stream commands and orders to control content of the printer buffer, referred to in 3270 subsystem context as Data Stream Compatible (DSC) or LU type 3, depending on the environment.
- IPDS -- Intelligent Printer Data Stream, a structured field data stream featuring all-points addressability (APA), for positioning of text, image,

and graphics on a page, and for control of media handling, duplexing, and downloading of fonts, symbol sets, overlays, and page segments.

Not all of them operate in both SNA and non-SNA environments. (See Figure 24.)

	SCS	IPDS	3270
SNA	LU type 1	LU type 1	LU type 3
Non-SNA	no	yes	DSC

Figure 24. Printer Data Stream Types

LU TYPE 1 OPERATION

The key to maximizing the throughput of medium and high speed printers is to maximize the overlap of the supply of data by the host through a 3X74 control unit and the print operation.

The LU type 1 protocol passes the data stream (SCS or IPDS) through the subsystem control unit, which is then processed (decoded) in the printer. (See also Chapter 4 on pass-through of long data streams.)

For LU type 1, larger RU sizes should be selected to minimize overhead processing. To maximize the overlap of printing with data transfer from the host (through the control unit) to the printer, specify pacing at "2" or more, and use "exception" rather than "definite" response. The amount of buffer space available with LU type 1 is determined by the printer adapter design, and therefore RU size and pacing N should be chosen so that:

Printer adapter buffer > (2N - 1) x RU size

If this condition is violated, buffer space in the control unit may be used for printer data stream storage. This may impair the performance of other devices attached to the control unit, and in a bad case, may even increase channel utilization.

Some printers include a compression and/or compaction function for use in LU type 1 operations. Both can reduce the length of data streams, depending on their content. This is especially important for the performance of remote subsystems. The use of these functions requires applicable support in the host.

Compression is useful when a data stream contains strings of the same character, for example, blanks. Compression reduces such strings to control characters with a character count, while decompression does the reverse.

Compaction reduces data stream length by packing two 'master' characters (of a limited, preselected set) into a single byte.

LU TYPE 3 AND DSC OPERATION

In the LU type 3 and DSC operations, the 3270 data stream is processed in the control unit, and the result is transferred into a screen-sized buffer in the printer adapter. (This may be set to be 1920, 2560, 3440, or 3564 characters, emulating model 2, 3, 4, or 5 screens, respectively.) The host should not send more data than can be held in this buffer, otherwise data will wrap, and data may be lost.

The host will send new data only after the printer acknowledges the successful printing of the previous buffer load. See the upper part of Figure 25 on page 94. Therefore, host and control unit activity do **not** overlap the printing operation. For high, and even medium-speed printers, this transfer/processing time during which no printing occurs may become significant compared to the actual printing time, and may cause sizable throughput reductions.

The reduction in average print speed can be demonstrated by the expression below. See Figure 25 on page 94 for explanation of the variables.

$$\text{LPM actual} = \frac{\text{Tpr}}{\text{Tcu1} + \text{Ttr} + \text{Tpr} + \text{Tcu2} + \text{Th}} \times \text{LPM max}$$

For example, for a 2000 LPM printer (30 milliseconds per line) with a 1920 character buffer that is printing a document with a 120 character average line width (allowing about 16 lines per buffer load), Tpr = 480 milliseconds. Assuming furthermore that Tcu1 = 100, Ttr = 120, Tcu2 = 5, and Th = 100, LPM actual equals:

$$\text{LPM actual} = \frac{480}{100 + 120 + 480 + 5 + 100} \times 2000 = 1193 \text{ LPM}$$

For a 200 LPM printer with the same assumptions, LPM actual would be 187 LPM, a much less drastic decrease in throughput.

These issues are also addressed in a Technical Bulletin entitled **Printers Attached to 327X Control Units - Basic Performance Concepts,** G320-5906-0.

EARLY PRINT COMPLETE FUNCTION (EPC)

The Early Print Complete function aims to introduce overlap between the transfer/processing phase and the printing operation. It is featured on most medium and high-speed printers that attach to 3X74 control units. EPC may be invoked in DSC and LU type 3 operations by entering this mode on the printer. The Set Printer Characteristics structured field (SPC) in Release 3 of 3174 Configuration Support-A allows an EPC function to be set and reset from the host, if supported by the printer.

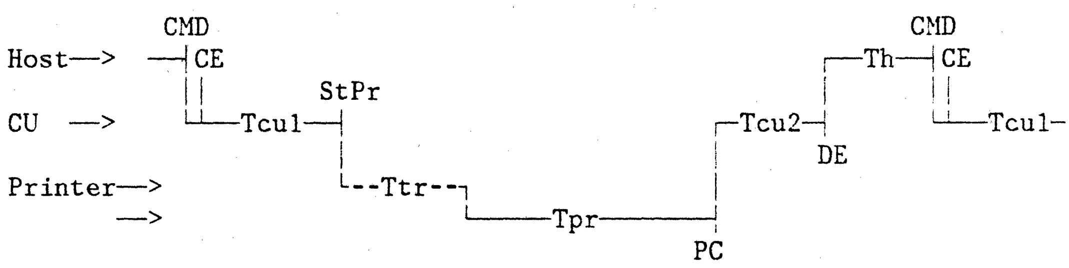

LU type 3 and DSC Operation

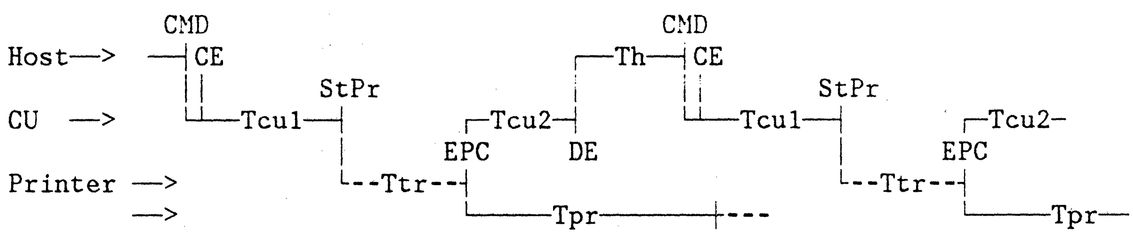

DSC (LU type 3) Operation with Early Print Complete

NOTES:

CMD	Command
CE	Channel end
DE	Device end
EPC	Early print complete
PC	Print complete
StPr	Start print command (on 3X74/4245 interface)
Tcu1	Control unit time (from CMD to Start-Print)
Tcu2	Control unit time (from Print-Complete to DE)
Th	Host response time (average)
Ttr	Printer buffer scan and transfer time (average)
Tpr	Print mechanism time

Figure 25. LU type 3 and DSC Operation, Early-Print-Complete Option

The EPC operational sequence is as follows. (See the lower part of Figure 25 on page 94.)

1. Host transfers data to the subsystem control unit.

2. Control unit processes data stream and fills printer adapter buffer.

3. Printer logic (backward) scans the (1920-byte) printer adapter buffer content, transfers data to subsequent buffer in printer, and signals host (through the control unit) that the printer adapter buffer is free to receive data again.

4. The data is processed, transferred to line buffers, and printed.

5. Before the printing mechanism has run out of lines to print, steps 1, 2, 3, and part of step 4 have readied new lines of data for printing.

6. Host fills request, and the sequence repeats.

Thus, the host and control unit times overlap partially or entirely with the actual printing operation.

Use of the EPC function may interfere with procedures for recovery from printer errors. When a printer error occurs in a block of data after its EPC has been sent, it is difficult for the host to establish whether the error occurred in that block or the next one.

When using EPC, the intervals to obtain the next buffer load of data and printing of the current load are partially overlapped. Using the numbers of our example:

$$Tpr > Tcu1 + Ttr + Tcu2 + Th$$

$$\text{because} \quad 480 > 325 \quad (= 100 + 120 + 5 + 100)$$

This suggests that there is total overlap, and the printer can run at maximum speed without interruption. It must be realized, however, that host and control unit times are averages, that is, they may sometimes combine to exceed Tpr in which case the printer runs out of data and temporarily stops anyway, which represents irretrievable printer throughput loss.

SUBSYSTEM CONTROL UNIT UTILIZATION BY PRINTERS

For a local subsystem, average control unit utilization as a result of printer operation is proportional with the rate at which it uses data. For a given data rate, utilization may differ somewhat with protocol.

For remote subsystem control units, the average utilization of the telecommunication link for a printer operation is proportional with the rate at which the printer uses data.

In band or wire-matrix printers, the electronics are designed to allow the printing mechanism to operate at maximum throughput capability, so that the latter usually determines the data rate requirements of the printer.

In page printers, when using large IPDS data streams for defining page content, the rate at which those data streams are processed in the printer may govern the page printing rate rather than the physical printing itself. Therefore, average data rates (in Kb per second) to be passed through the control unit may not increase as much as multiplying data steam length per page (in Kb) with the nominal page rate of the printer (in pages per second) would suggest.

PRINTER DATA STREAM BENCHMARKS

Printer performance data for DSC and LU types 1 and 3 in these guidelines were obtained with so-called ripple benchmarks and sometimes with same-character benchmarks. Ripple benchmarks are generated by repeating the same character sequence to complete a line, with a small displacement in consecutive,

single-spaced lines. Each line ends with an NL order. With a same-character benchmark, only a single character is used rather than a sequence of characters. With some printers, print speed may be sensitive to the character content of the data stream.

Each of four benchmarks have lines of identical lengths, specifically 40, 70, 100, and 130 characters long.

Although these data streams are unlikely to show up in actual applications, they do permit more accurate assessment of host/control unit performance because they avoid the introduction of hard-to-quantify variations encountered in actual applications.

Averaging the widths of nonblank lines in an application data stream allows assessment of its performance by comparing it with the performance of a benchmark of comparable width.

CHAPTER 12. IBM 3174 WITH ASYNCHRONOUS EMULATION ADAPTER

This chapter provides performance information on IBM 3174 Models 1L, 1R, 2R, and 3R Subsystem Control Units with up to three Asynchronous Emulation Adapters (AEA, #3020), each capable of attaching a combination of up to eight ASCII display stations, printers, and ASCII host systems. Models 51R, 52R, and 53R can accommodate one adapter. (See Figure 26.)

The information in this chapter supersedes an earlier version of this data, which was made available on HONE in FLASH 8738.

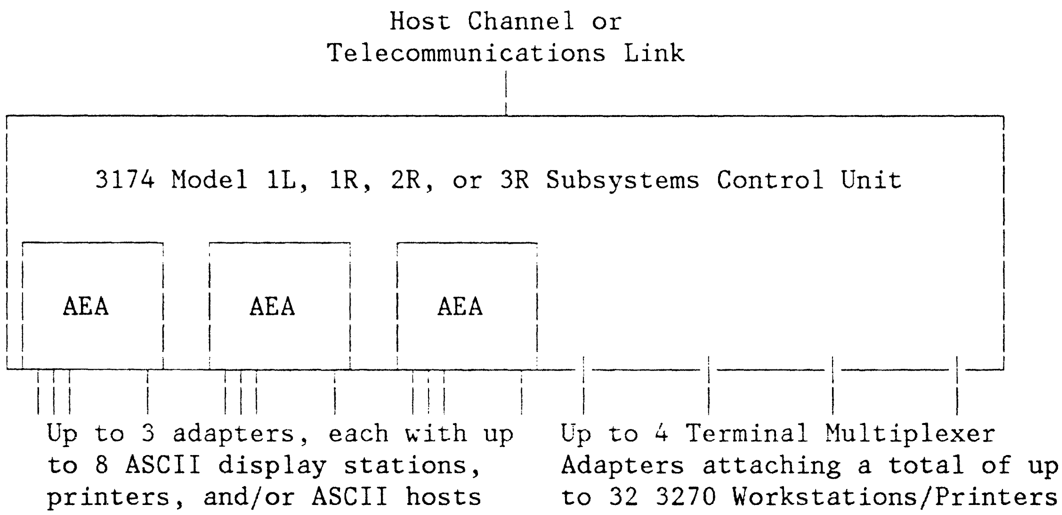

Figure 26. IBM 3174 Control Unit with Asynchronous Emulation Adapters

Each AEA port provides an EIA RS-232-C electrical interface, and supports asynchronous transmission speeds of 300, 600, 1200, 2400, 4800, 9600, and 19200 bps via modems over switched or nonswitched communications facilities, or via direct connection, without modems.

The asynchronous emulation adapter supports three modes of operation:

- **ASCII Terminal Emulation:** IBM 3270 displays or printers emulating an IBM 3101 or other ASCII display, or ASCII printer, for connection to ASCII hosts or public data networks. (Sometimes referred to internally as RPC: reverse protocol conversion.)

- **3270 Terminal Emulation:** ASCII displays or printers emulating an IBM 3178 Model C20, 3279 Model 2A, or 3287 Model 2, for connection to an IBM host. (Sometimes referred to as FPC: forward protocol conversion.)

- **ASCII Pass-Through:** ASCII terminals connecting to ASCII hosts or public data networks through the 3174. (Sometimes referred to as APT: ASCII pass-through.)

Only the 3174 subsystem performance is discussed here; host (IBM or ASCII) and network (IBM or ASCII) performance information is not included.

The maximum number of sessions that can be addressed in a non-SNA local or remote 3174 is 32 (180 in the SNA version). Therefore, using AEA ports in a non-SNA 3174 limits the number of addresses available for 3270 terminals.

ASCII TERMINAL EMULATION WITH A 3270 DISPLAY STATION

IBM 3270 display stations attached to a 3174 port, and operating in CUT mode, may be used for emulation of an IBM 3101 or other ASCII terminal. Workstations emulating a 3270 display station may be used for ASCII terminal emulation but cannot transfer files when operating in this mode.

BENCHMARK DATA STREAMS

The 3270 benchmarks specified in Appendix A do not apply to ASCII hosts.

For characterization of the performance of the **scrolling mode** of operation, a 5760-character ASCII benchmark (seventy-two 80-character rows) for outbound transmission was adopted. Characters start being displayed at the top of the screen and progress downward until all rows are filled. Subsequent characters will cause an upward scrolling movement of one or more rows at a time, depending on the rate they are received from the ASCII host and other activity in the 3174 and the emulation adapter. Upon completion of the transmission, the last 24 rows of the benchmark are displayed.

Panel mode implies that the outbound data is constrained to the physical screen boundaries, and that no scrolling takes place. The host paints the screen in 3270 fashion by issuing "move cursor" commands, etc. No benchmarks for ASCII host panel mode applications have been adopted. See the section on Panel Mode Performance in this chapter for analysis of a sample application which may be used as a guide for analysis of specific applications.

INBOUND TO ASCII HOST

Keystrokes from the 3270 workstation are sent to the ASCII host for processing. After processing, the ASCII host echoes changes to displayed characters or the cursor position back to the workstation. Keystroke content can vary widely and should be estimated for each application.

The 3174 keystroke processing time is negligible. The emulation adapter processing time is about 10 milliseconds for character and cursor move keystrokes, and about 100 milliseconds for insert, delete and erase EOF. These times do not include ASCII host processing or the time on the ASCII link to and from the ASCII host.

In most cases, keystrokes will be processed without noticeable delay. However, updating of the screen for typematic insert may lag behind the keying rate because the 3270 terminal typematic rate is 10 cps (100 milliseconds per keystroke), while the sum of the delays in the emulation adapter, ASCII host, and ASCII link will, in some cases, exceed 100 milliseconds. Therefore, insert keystrokes in typematic mode may be accumulated in the emulation adapter, and appear in bursts of more than one character moved on the screen.

SCROLLING MODE PERFORMANCE

The response time of the benchmark for scrolling is defined as the time from the start of its transmission over the ASCII host link until completion of screen update.

The emulation adapter limits the rate at which data is received by means of XON/XOFF flow control to a maximum of 575 characters per second for a single terminal for the purpose of producing a smooth scrolling action. (The entire screen of a 3270 display is updated for each scrolling action. When the ASCII transmission rate is too high, the screen update will contain multiple new rows.)

When multiple terminals are operating simultaneously on a single adapter, then the aggregate rate increases to a maximum of 610 characters per second. Therefore, for ASCII link rates of 9600 bps or higher, response time is independent of the ASCII rate.

Measurement data for a single emulation adapter is shown in Figure 27.

	# of Terminals Operating 1	2	3	4
Response time, seconds	10*	20	30	40
Characters per second	575	590	610	610
AEA utilization, percent	84	94	95	97
3174 micro utilization, percent	14	**	**	22

NOTES: * 12 for 4800 bps; ** Not measured

- Benchmark is seventy-two 80-character rows
- ASCII host link rates 4800 to 19,200 bps.

Figure 27. Performance of ASCII Terminal Emulation

For ASCII rates of 2400 bps or less, the maximum aggregate character rate is proportional to the rate:

Characters per second = ASCII bps rate / 10 bits per character

Each terminal will receive an equal distribution of the character rate.

For other file sizes, the response time will be approximately proportional to the length.

PANEL MODE PERFORMANCE EXAMPLE

The following example is used to serve as a guide:

- Three AEAs; four active terminals per AEA
- One transaction per minute per active terminal.
- Per transaction: 200 characters (keystrokes) inbound to ASCII host, 1000 characters outbound from ASCII host.
- ASCII rate = 9600 bps.

As a general rule, the utilization of both the emulation adapter and the 3174 microprocessor should not exceed 65 percent.

The **outbound response time** for a non-overlapping transaction is (values below are from Figure 27 on page 100):

    Response time = 1000 / 575 = 1.7 seconds

When they all overlap, then:

    Response time = 1000 / (610/4) = 6.6 seconds

The probability of overlap is low, since:

    4 x 1.7 = 6.8 seconds (11 percent of 60 seconds)

Therefore, the outbound response time should be close to 1.7 seconds.

The **AEA outbound utilization** is figured by noting that each transaction produces a utilization of 84 percent during the 1.7 second transaction response time, therefore:

    AEA service time = 0.84 x 1.7 = 1.4 seconds per transaction

Because utilization equals transaction rate times service time:

    AEA utilization = 4 x (1/60) x 1.4 = 9.3 percent

The **inbound response time** depends on keystroke response time. (See section "Inbound to ASCII Host" in this chapter.) Unless the adapter is over-utilized, the keystroke processing will keep up with user keystrokes.

The **AEA inbound** utilization is:

    AEA service time = 200 x .010 = 2.0 seconds per transaction,

and

    AEA utilization = 4 x (1/60) x 2 = 13.3 percent,

and therefore:

    Total AEA utilization = 9.3 + 13.3 = 22.6 percent

Since the emulation adapter utilization is much less than 65 percent, its performance in this application should be satisfactory if the response time meets the objectives.

With respect to **3174 performance**, the following comments apply:

- **Outbound response time** – Unless the 3174 utilization is high, it will keep up with the emulation adapter data rate. If the utilization is high, about 65 percent or more, it may fall behind. However, any additional delay is unlikely to be perceptible due to the much faster rate at which the 3174 microprocessor processes data relative to the emulation adapter rate.

- **Outbound utilization** – It was determined above that there is little probability of overlap of transactions for one emulation adapter. Therefore the 3174 microprocessor utilization will be about 14 percent (from Figure 27 on page 100) for the duration of each outbound transaction.

  Service Time = 0.14 x 1.7 = 0.24 second per transaction
  Utilization = 3 x 4 x (1/60) x 0.24 = 4.8 percent

- **Inbound response time and utilization** – The inbound data consists of keystrokes which have negligible response times and utilizations.

The 3174 utilization appears to be very low so it should not be a factor in determining overall performance of this application.

EFFECT OF ASCII TERMINAL EMULATION ON OTHER 3270 DISPLAY STATIONS

ASCII terminal emulation uses the same 3174 priority levels as MFI operations by other 3270 workstations. However, 3174 utilization by the emulation per comparable 3270 transaction is much larger in both scrolling and panel mode. (A single CUT mode 3270 workstation processing one type A-1200 transactions per minute utilizes an average of about 0.16 percent of a 3174.)

One can expect, therefore, the performance of 3270 workstations interacting with an IBM host to be significantly degraded. For example, when three 3270 terminals, each connected to an ASCII host through a separate emulation adapter, are scrolling simultaneously, 3174 utilization will be 42 percent, that is, two thirds of the recommended maximum of 65 percent.

3270 TERMINAL EMULATION WITH AN ASCII DISPLAY STATION

When emulating a 3270 terminal function with an AEA-attached ASCII display station, outbound data streams are processed in the 3174, and changes in the appearance of the screen as specified by the 3174 are asynchronously transmitted by the emulation adapter to the ASCII terminal at the speed of the line connecting that terminal.

This requires that the ASCII terminal be operated in echo mode, that is, data entered from the keyboard is sent inbound only, while its screen is only affected by commands and data received in outbound data streams.

Because in these operations an IBM host uses 3270 data streams, types A-1200 and C-1200 benchmarks were selected for the presentation of performance data. (See Appendix A.)

RESPONSE TIMES TO LAST AND FIRST CHARACTER

Subsystem response time to last character **does not** include host processing time and network delays. (See section "Response to Last Character" in Chapter 7.)

Response time to first character has wide usage with remotely attached terminals because slow link transmission rates usually result in long response times to last character. However, if the host and network delays are not included, as is the case here, then the time to first character is almost immediate for 3270 emulation with an ASCII terminal. The time to first character is very short for the following reasons:

- For local channel host attachment, the 3174 and the emulation adapter start processing the outbound data as soon as the first SNA RU or non-SNA block has been received from the channel. The transmission time at channel rates is insignificant, and the time between end-of-channel transmission and the start of ASCII port transmission of characters will be typically less than 100 milliseconds.

- For remote host attachment, the 3174 and the emulation adapter start processing the outbound data as soon as the first SDLC segment or BSC block has been received from the network. The time between the end-of-network transmission of the first segment or block and the start of ASCII port transmission is again typically less than 100 milliseconds.

- The ASCII port transmission speed has no significant effect on the inbound delay because a copy of the ASCII display is kept current in the emulation adapter. Therefore, when the Enter key is depressed, there is no need to move the ASCII display buffer at the relatively slow ASCII port speed to the 3174/emulation adapter for inbound processing. The time required for inbound processing, even if the entire display buffer is sent inbound, is again less than 100 milliseconds.

Therefore, the 3174/emulation adapter contribution to response time to first character is less than 200 milliseconds.

LOCAL 3174/AEA/ASCII TERMINAL RESPONSE TIMES

The table in Figure 28 on page 104 may be used to estimate, for a local 3174/AEA subsystem, the response times to last displayed character. The times apply to both SNA and non-SNA environments. "A" designates the type A-1200 benchmark, and "A/C" the alternate transmission of A-1200 and C-1200 benchmarks. (The "C" type data stream modifies an existing panel, in this case the "A" panel.)

	AEA Port Speed, bps									
	1200		2400		4800		9600		19,200	
CTR	A	A/C	A	A/C	A	A/C	A	A/C	A	A/C
1	12.7	10.6	6.5	5.5	3.4	2.9	2.0	1.7	1.4	1.1
5	12.7	10.6	6.5	5.5	3.4	2.9	2.1	1.7	1.5	1.2
10	12.7	10.6	6.5	5.5	3.4	2.9	2.1	1.7	1.5	1.2
15	12.7	10.6	6.5	5.5	3.5	2.9	2.1	1.7	1.6	1.2
20	12.9	10.7	6.5	5.5	3.5	2.9	2.2	1.7	1.7	1.2
25	13.4	10.9	6.5	5.5	3.6	2.9	2.3	1.7	1.9	1.2
30		11.3	6.6	5.5	3.7	2.9	2.4	1.8	2.1	1.3
35			6.8	5.5	4.0	2.9	2.6	1.8	2.3	1.3
40			7.0	5.5	4.4	3.0	2.9	1.9	2.7	1.4
45				5.5		3.0	3.3	1.9	3.1	1.4
50				5.6		3.1		2.0		1.6
55				5.7		3.2		2.1		1.6
60				5.9		3.3		2.3		1.7
65						3.4		2.5		1.9
70										2.0

NOTES:

- Response times in seconds, not provided for AEA utilizations in excess of 65 percent
- CTR: transaction rate, per AEA per minute
- At 1.2 kbps, CTR is limited by the Port Speed

Figure 28. 3270 Emulation with ASCII Terminal on Local 3174/AEA

The table has been truncated for values that cause the emulation adapter utilization to exceed 65 percent. An AEA can operate at 100 percent utilization, but a maximum of 65 percent is recommended to avoid excessive queuing delays and variations in response times. The table should be used in the following way:

1. Estimate the number of transactions per minute per active ASCII terminal.

2. Multiply by the number of active ASCII terminals to get the total control unit transaction rate (CTR) associated with all AEA devices. Note that physically attached terminals not in use, even if logged on, are not considered active. For many applications, it is not unusual for only 25 percent, or less, to be active.

3. If there is more than one emulation adapter, then allocate the total CTR among the adapters. If there is a different CTR for each adapter, then there will be different response times associated with each adapter.

4. Choose either the "A", or the "A/C" column as the most closely approximating the application.

5. Read the response time from the "A" or "A/C" columns for the ASCII port speed in kbps and CTR for each emulation adapter.

6. Interpolate between response time points in the table, if it does not include values at the application operating point.

7. The emulation adapter assigns a service priority to each ASCII port. The highest port number has the highest service priority. Experiments with the data streams in Figure 28 on page 104 indicate that this service priority does not affect response times significantly, in most cases.

   The response time of terminals attached to the highest port number may be fifteen percent less than the values in the table, and those attached to the lowest port number may be fifteen percent more. This effect only occurs when transactions are queued up, and there is contention between ports for service by the AEA microprocessor.

   However, it has been found that this difference can be much higher (100 percent or more) for certain environments. They are categorized by the type and data rate of the host connection, the ASCII port speed, and the number of active terminals. The critical environments are shown in Figure 29 on page 106. When both the number of active terminals and the CTR rates are equal to, or greater than the values in the table, response time differences of more than 30 percent can occur. Since the CTR rates are relatively high, they may occur only infrequently, and only during peak periods of the day when more terminals are expected to be active. Even if the average CTR rate is predicted to be less than the critical value, there can be short periods when the low priority ports will experience long delays.

   For example, when the Enter key is depressed simultaneously on three or more terminals running the "A" data stream with local host attachment and a 19.2 kbps ASCII rate, then terminals attached to ports other than the two high priority ones will experience much longer response times. This situation could also occur when the two highest priority port terminals are paging through a file as rapidly as possible.

   Production type applications may be most sensitive to this effect. All attached terminals may be active, and lower productivity with some terminals could result in user complaints.

   Users in professional type environments, where twenty percent or less of the attached terminals may be active, may encounter occasional long response times that do not impact the user significantly; users would not find this effect objectionable.

		AEA Port Speed, bps					
		4800		9600		19,200	
Data Stream	Host Attachment	CTR	Active Term's	CTR	Active Term's	CTR	Active Term's
A	Local	50	5	50	4	50	3
	19,200 bps	45	5	45	4	45	3
	9600 bps	–	–	45	7	45	5
A/C	Local	70	5	70	3	70	3
	19,200 bps	55	5	55	3	55	3
	9600 bps	55	6	55	5	55	5

NOTE: CTR transaction rate, per AEA per minute (– for AEA utilizations in excess of 65 percent)

Figure 29. 3270 Emulation with ASCII Terminals - Critical Environments

The response times for low CTR values in Figure 28 on page 104 may be approximated by the expressions:

For "A": Response time = 0.6 + (14500 / line speed) seconds
For "A/C": Response time = 0.5 + (12100 / line speed) seconds

They may be used for estimating response times for other line speeds than those shown.

RESPONSE TIMES FOR OTHER CASES

The following characteristics of this mode can be used to predict performance for cases other than those presented here:

- Inbound response time is very small, and can be ignored regardless of data stream length or content. See section "Response Times To Last and First Character" in this chapter.
- Response time is approximately proportional to outbound data stream length for a particular type of data stream.
- The maximum recommended CTR at an emulation adapter utilization of 65 percent will be approximately inversely proportional to outbound data stream length for a particular type of data stream.
- Response time is approximately inversely proportional to the transmission rate of the AEA port for speeds up to 9600 bps. Further increases in speed up to 19,200 bps may not result in a proportional reduction in response time due to emulation adapter processing capability.

- The capacity of the 3174 in terms of total transaction rate is directly proportional to the number of emulation adapters. This is because almost all of the processing is done by the emulation adapter. The 3174 microprocessor is utilized for about 100 milliseconds per transaction and cannot normally be stressed by peak emulation adapter transaction rates. An exception would be for applications with very short data lengths, such as one-line writes to operator consoles. However, the total CTR would have to exceed 400 per minute before queuing delays in the 3174 processor became significant.

- The ASCII transmission protocol for these response times is 10 bits per character. This is common to almost all applications, although a few may use 11. In this case the response times will increase proportionally by 10 percent.

RESPONSE TIMES OF REMOTE 3174S

As explained above, the subsystem response time is very nearly a function of outbound processing time. Almost immediately after an SDLC segment is received from the network, the 3174 transfers it to the emulation adapter for processing. Processing of a BSC block starts as soon as its first ten characters have arrived. Emulation adapter translation to ASCII and transmission of the new image to the ASCII terminal then proceeds overlapped with network transmission of any succeeding segments or block data. The schematic in Figure 30 illustrates this process.

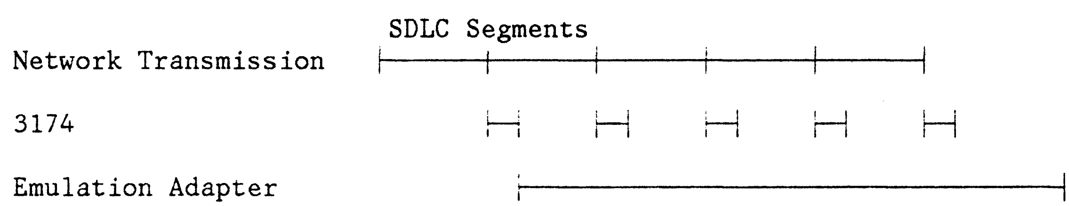

Figure 30. Processing/Transmission Overlap

Therefore, the response time will be approximately the larger of either the network transmission time or the emulation adapter time.

There may be queuing or other delays associated with the data passing between network nodes. These delays are not accounted for here, and should be determined by using available network analysis tools such as FIVE or SNAPSHOT.

The following possibilities exist:

- **AEA port ≤ network** – When AEA port speed is smaller than, or equal to network speed, the response times will be approximately equal to the greater of the times in Figure 28 on page 104, or of those predicted by network performance analysis. For low transaction rates, the emulation adapter will be the bottleneck, just as it is for local channel attachments. Follow the instructions for determining local channel-attached response times. Due

to the internal structure of the emulation adapter, its response time may not increase as rapidly as that of a telecommunications link with increasing utilization. Therefore, network response times may become the dominant factor at higher transaction rates.

- **AEA port > network** – When AEA port speed is greater than network speed, the response times should be determined by network performance analysis since the network is the bottleneck.

- **More than one emulation adapter** – The communications link will carry the entire transaction load, and each emulation adapter only a portion of it. Response times for each emulation adapter should be determined from the local channel table above. The larger of the network and emulation adapter response times should be used.

EFFECT OF ASCII AND 3270 WORKSTATIONS ON EACH OTHER

In local 3174s, 3270 emulation with ASCII terminals can be affected by heavy MFI transaction loads on 3270 workstations because some of the functions performed by the 3174 microprocessor for the emulation adapter(s) have a lower priority level than 3270 support. Since 3270 functions are usually executed within 100 milliseconds, this interference should not be noticeable except at unusually high 3270 transaction rates, say 300 per minute, or more. The increase in response time will approximate those indicated in this guide for 3270 CUT-mode workstations.

To assess the effects of MFI transactions at ASCII terminals on 3270/host operations, AEA transactions may be added to 3270 terminal transactions for the purpose of determining their response times. This will produce conservative results since 3270 transactions get priority when contending with the emulation adapter for the 3174 microprocessor.

In remote 3174s, both 3270 workstation and AEA transactions, as applicable, should be included in network performance analysis.

ASCII PASS-THROUGH

In an ASCII pass-through operation, an ASCII device communicates (in character mode) with an ASCII host through an emulation adapter. When the host and terminal links have different speeds, performance is governed by the link with the lowest speed.

A data stream of any length is simply passed through from one AEA port, operating in receive mode, to another port in the same or different AEA, operating in transmit mode.

AGGREGATE RATES

Each emulation adapter can support a finite maximum aggregate data rate: transmit, receive, or a combination of the two. This aggregate rate is the sum of the rates of ports that are operating simultaneously. For transmit, the maximum aggregate rate is 24,000 bps and for receive it is 55,000 bps. If there is a mix of transmit and receive then the maximum aggregate rates of each should fit the following equation:

$$1 \geq (\text{Transmit Rate}/24{,}000) + (\text{Receive Rate}/55{,}000)$$

It should be understood that that this is not a limitation on the number of physical ports or their speed. The aggregate data rate limit applies to the number and speed of the ports that may be in use simultaneously.

If XON/XOFF flow control is in use, then transmit and receive will be interrupted as required to constrain the rates within the above limits. However, if the ASCII sources attempt to exceed the limit, then there can be short periods when receive characters are lost. This occurs because flow control may not be invoked fast enough to prevent overrun of the ASCII port interface logic. Therefore, applications with heavy receive traffic should be evaluated carefully to assure that the aggregate data rates are not exceeded if occasional lost characters will be a problem.

If flow control is not in use then buffer overrun errors can occur.

The utilization of the 3174 microprocessor by this mode is small and can generally be ignored.

APPENDIX A - PERFORMANCE DATA SHEETS FOR 3X74-ATTACHED DEVICES

This appendix contains Performance Data Sheets for IBM display stations, workstations, and printers attaching to IBM 3174 and 3274 Control Units, filed by product number.

During the life of this edition, new or amended Performance Data Sheets may become available on the HONE System.

PERFORMANCE DATA SHEET FOR IBM 3180 AND 3192 C/D DISPLAY STATIONS

The performance characteristics of the IBM 3192 Models C and D are about the same as those of the IBM 3180 display stations. They both operate in CUT mode exclusively.

The table in Figure 8 on page 47 lists single-thread response times and control unit utilization data for the 3180 display station (for several screen sizes).

The IBM 3180/3X74 subsystem response characteristics differ in a minor way from those of a 3278/3X74 subsystem, because of the manner in which 3180 display buffer data are displayed on the screen. The response data in the table shows the return of the host acknowledgment to be faster for the 3180 than for the 3278. However, where the 3278 displays the last character virtually simultanegusly with the return of this signal, there is a small delay in a 3180 (and 3192 models C and D) averaging about 35 millisecond for a model 2 screen, 45 for a model 4, and 75 milliseconds when a 96x80 character scroll buffer is used. This buffer transfer delay in the 3180 display station is independent of control unit utilization.

The result is that, compared with the 3278, the response-to-last-character of these display stations on lightly loaded control units can be up to twenty percent more, with the difference becoming less with increasing control unit load.

Processing a data stream in an explicit partition yields about the same 3180/3274 response time as in an implicit partition of the same size. However, when the size of the explicit partition was quadrupled, the benchmarks with an Erase/Write command exhibited a substantial increase (25-35 percent) in subsystem response, due to the time required to erase the partition.

The 3180, when operating with a large partition size (>4 Kb) and a substantially smaller outbound data stream, may exhibit better performance with the 3274-41 than with the 3174 Model 1L.

PERFORMANCE DATA SHEET FOR IBM 3X74/3193 IMAGE DISPLAY STATION

The IBM 3193 Display Station can display both image and alphameric data. It operates with the 3X74 control units in DFT mode. This data sheet describes image performance only. Its alphameric performance characteristics are similar to those of other DFT devices. (See Chapter 6.)

OPERATIONAL DESCRIPTION

The 3193 resolution is 100 (horizontal) by 100 (vertical) PELs per inch (ppi). Source documents scanned at this resolution can be mapped directly onto the 3193 screen in their entirety. If they are scanned at higher resolutions then there are two possibilities: - the entire source document can be displayed but with PELs dropped by the 3193 to make it fit the physical screen, or a part of the source document can be displayed with all its PELs intact. A Set Resolution Modification Mode instruction in the host data stream is used to select one of these modes. For the data given here, the mode option is "on", displaying the entire source document with excess PELs dropped by the 3193.

The image data is first decompressed by the 3193, and then painted on the screen. Orders can be included in the host data stream which manipulate the decompressed image in various ways such as scale, rotate, mirror and invert, prior to painting it on the screen. The response times in this data sheet are for the simple display of the source image, that is, without manipulation. Additional time for image manipulation may be added.

BENCHMARKS FOR IMAGE DISPLAYS

The performance of two A-4 size CCITT (International Telegraph and Telephone Consultative Committee) benchmark images, each scanned in three resolutions, is described here. The resolutions are 200(H) by 100(V), 200(H) by 200(V), and 240(H) by 240(V) ppi. The CCITT image benchmark "Business Letter" (No 1, Figure 32 on page 118) represents a relatively sparse image. CCITT benchmark "French Text" (No. 4, Figure 33 on page 119) represents a denser image calling for more data. See Figure 31 on page 116 for the data stream size associated with each.

RESPONSE TIMES WITH IMAGE BENCHMARKS

The performance of these benchmarks is summarized in Figure 31 on page 116. If each 3193 is requesting images at a rate of N per minute, then the number of 3193s that can be supported is CTRmax/N. The 3274-41A values assume that RPQ #8K1311, DFT Display Performance Enhance, has been installed. The function of this RPQ is included in the 3174-1L microcode base.

Image Benchmark --->		Business Letter			French Text		
Horizontal, ppi		200	200	240	200	200	240
Vertical, ppi		100	200	240	100	200	240
Size, kbytes		12.3	18.4	21.9	46.5	70.2	81.9
3274-41D[3]	RT-0%	5.6	6.0	6.5	9.4	11.8	13.3
	RT-65%	5.7	6.4	7.0	10.4	13.0	14.6
	CTRmax	190	126	107	51	35	29
3274-41A[1]	RT-0%	5.3	5.9	6.4	9.1	11.1	12.5
	RT-65%	5.6	6.4	7.0	10.9	14.4	15.8
	CTRmax	85	58	48	24	16	14
3174-1L[3] (non-SNA)	RT-0%	5.5	6.0	6.5	9.3	11.7	13.2
	RT-65%	5.6	6.3	6.9	10.1	12.6	14.2
	CTRmax	221	148	124	60	41	34
3174-1L[2] (SNA)	RT-0%	5.3	5.8	6.3	8.9	10.7	12.2
	RT-65%	5.4	6.0	6.5	9.3	11.4	13.2
	CTRmax	195	139	110	56	38	32

NOTES:

RT-0% Time (seconds) from Enter/PF key to completion of image decompression and display by the 3193 with no other 3X74 utilization; host delay = 0.

RT-65% As RT-0%, but with 65 percent 3X74 utilization.

CTRmax CTR (image benchmarks per minute) at maximum recommended 3X74 utilization of 65 percent.

[1] RU size = 1536 bytes, pacing = 2
RPQ 8k1311 installed

[2] RU size = 2048 bytes, pacing = 2

[3] Block size = 7168 bytes

Figure 31. IBM 3193 Display Station Performance (Image Display)

There is little difference between RT-0% and RT-65%, because the decompression and painting of the image in the 3193 accounts for most of the response time. The increase in response time between RT-0% and RT-65% is due to queuing delays in the 3X74, but these are small compared to the 3193 time.

The 3193 processing time will increase when image manipulation orders are included in the host data stream. The 3X74 time remains the same, because the additional amount of data passed through is not significant. Add about one second to the response times in Figure 31 for image editing that extracts four sections from the source image, manipulates each of these with orders such as

Scale, Rotate, Mirror and Invert, and then places the sections at specified locations.

RESPONSE TIMES FOR OTHER IMAGES

For displaying a full A-4 size image on the screen, the 3193 processing time may be approximated with the following equation:

$$T = (Kb/DC) + P \quad \text{second}$$

where:

DC = Decompress rate = 11.2 kbytes per second
Kb = Image size, in kbytes
P = Paint screen = 4.1 second (without image manipulation orders).

Note: For microcode release D4.2/N4.2 or higher, one may expect the paint time P to be several seconds less than indicated. The same is true for the "RT-0%" and "RT-65%" values in Figure 31 on page 116.

The 3X74 component of response time is proportional to the image data length and is small compared to the 3193 time. Approximate total 3X74/3193 response time can be computed by adding the following percentages to the 3193 time calculated with the above equation.

	SNA	Non-SNA
RT-0%	10%	15%
RT-65%	15%	22%

The CTR at maximum recommended 3X74 utilization of 65 percent is inversely proportional to the image data length. Therefore, CTRs for other images can be computed by relating their data lengths to those in Figure 31 on page 116.

Data lengths for other source document scanning resolutions are related to the number of PELs. For a hundred percent increase in PELs the data length increases about fifty percent.

For the effect of the often lengthy image data streams on MFI performance of CUT workstations on the same control unit, refer to Chapter 4.

THE SLEREXE COMPANY LIMITED

SAPORS LANE - BOOLE - DORSET - BH 25 8 ER

TELEPHONE BOOLE (945 13) 51617 - TELEX 123456

Our Ref. 350/PJC/EAC 18th January, 1972.

Dr. P.N. Cundall,
Mining Surveys Ltd.,
Holroyd Road,
Reading,
Berks.

Dear Pete,

Permit me to introduce you to the facility of facsimile transmission.

In facsimile a photocell is caused to perform a raster scan over the subject copy. The variations of print density on the document cause the photocell to generate an analogous electrical video signal. This signal is used to modulate a carrier, which is transmitted to a remote destination over a radio or cable communications link.

At the remote terminal, demodulation reconstructs the video signal, which is used to modulate the density of print produced by a printing device. This device is scanning in a raster scan synchronised with that at the transmitting terminal. As a result, a facsimile copy of the subject document is produced.

Probably you have uses for this facility in your organisation.

Yours sincerely,

Phil.

P.J. CROSS
Group Leader - Facsimile Research

(at 74 percent of actual size)

Figure 32. CCITT Image Benchmark (No. 1) "Business Letter"

L'ordre de lancement et de réalisation des applications fait l'objet de décisions au plus haut niveau de la Direction Générale des Télécommunications. Il n'est certes pas question de construire ce système intégré "en bloc" mais bien au contraire de procéder par étapes, par paliers successifs. Certaines applications, dont la rentabilité ne pourra être assurée, ne seront pas entreprises. Actuellement, sur trente applications qui ont pu être globalement définies, six en sont au stade de l'exploitation, six autres se sont vu donner la priorité pour leur réalisation.
Chaque application est confiée à un "chef de projet", responsable successivement de sa conception, de son analyse-programmation et de sa mise en oeuvre dans une région-pilote. La généralisation ultérieure de l'application réalisée dans cette région-pilote dépend des résultats obtenus et fait l'objet d'une décision de la Direction Générale. Néanmoins, le chef de projet doit dès le départ considérer que son activité a une vocation nationale donc refuser tout particularisme régional. Il est aidé d'une équipe d'analystes-programmeurs et entouré d'un "groupe de conception" chargé de rédiger le document de "définition des objectifs globaux" puis le "cahier des charges" de l'application, qui sont adressés pour avis à tous les services utilisateurs potentiels et aux chefs de projet des autres applications. Le groupe de conception comprend 6 à 10 personnes représentant les services les plus divers concernés par le projet, et comporte obligatoirement un bon analyste attaché à l'application.

II - L'IMPLANTATION GEOGRAPHIQUE D'UN RESEAU INFORMATIQUE PERFORMANT

L'organisation de l'entreprise française des télécommunications repose sur l'existence de 20 régions. Des calculateurs ont été implantés dans le passé au moins dans toutes les plus importantes. On trouve ainsi des machines Bull Gamma 30 à Lyon et Marseille, des GE 425 à Lille, Bordeaux, Toulouse et Montpellier, un GE 437 à Massy, enfin quelques machines Bull 300 TI à programmes câblés étaient récemment ou sont encore en service dans les régions de Nancy, Nantes, Limoges, Poitiers et Rouen ; ce parc est essentiellement utilisé pour la comptabilité téléphonique.
A l'avenir, si la plupart des fichiers nécessaires aux applications décrites plus haut peuvent être gérés en temps différé, un certain nombre d'entre eux devront nécessairement être accessibles, voire mis à jour en temps réel : parmi ces derniers le fichier commercial des abonnés, le fichier des renseignements, le fichier des circuits, le fichier technique des abonnés contiendront des quantités considérables d'informations.
Le volume total de caractères à gérer en phase finale sur un ordinateur ayant en charge quelques 500 000 abonnés a été estimé à un milliard de caractères au moins. Au moins le tiers des données seront concernées par des traitements en temps réel.
Aucun des calculateurs énumérés plus haut ne permettait d'envisager de tels traitements.
L'intégration progressive de toutes les applications suppose la création d'un support commun pour toutes les informations, une véritable "Banque de données", répartie sur des moyens de traitement nationaux et régionaux, et qui devra rester alimentée, mise à jour en permanence, à partir de la base de l'entreprise, c'est-à-dire les chantiers, les magasins, les guichets des services d'abonnement, les services de personnel etc.
L'étude des différents fichiers à constituer a donc permis de définir les principales caractéristiques du réseau d'ordinateurs nouveaux à mettre en place pour aborder la réalisation du système informatif. L'obligation de faire appel à des ordinateurs de troisième génération, très puissants et dotés de volumineuses mémoires de masse, a conduit à en réduire substantiellement le nombre.
L'implantation de sept centres de calcul interrégionaux constituera un compromis entre : d'une part le désir de réduire le coût économique de l'ensemble, de faciliter la coordination des équipes d'informaticiens; et d'autre part le refus de créer des centres trop importants difficiles à gérer et à diriger, et posant des problèmes délicats de sécurité. Le regroupement des traitements relatifs à plusieurs régions sur chacun de ces sept centres permettra de leur donner une taille relativement homogène. Chaque centre "gèrera" environ un million d'abonnés à la fin du VIème Plan.
La mise en place de ces centres a débuté au début de l'année 1971 : un ordinateur IRIS 50 de la Compagnie Internationale pour l'Informatique a été installé à Toulouse en février ; la même machine vient d'être mise en service au centre de calcul interrégional de Bordeaux.

Photo n° 1 - Document très dense lettre 1,5mm de haut -
Restitution photo n° 9

Figure 33. CCITT Image Benchmark (No. 4) "French Text"

PERFORMANCE DATA SHEET FOR IBM 3194 DISPLAY STATION

This performance data sheet provides subsystem response times and utilizations, and file transfer data for IBM 3194 Display Stations attached to a local IBM 3174 Model 1L or IBM 3274 Model 41 control unit.

Refer to Chapters 6 and 8 for an introduction to the use of the response time and file transfer information, respectively.

The IBM 3194 Display Station can be customized for operation in CUT mode or DFT mode. (DFT mode on 3274 only with Configuration Support D release level 64.2 or later.)

In CUT mode, the 3194 operates as a single-screen display station with 1920 characters (24 rows of 80 characters). No 7-color or extended highlighting function is available in this mode.

In DFT mode, 3194 Models D and H (except H50) can emulate display stations with model 2, 3, 4, and 5 screens. (Models C and H50 can emulate model 2 and 3 screens.) Up to four host sessions can be conducted simultaneously in DFT mode, provided the 3X74 control unit port is customized for the appropriate number of sessions.

SUBSYSTEM UTILIZATION AND RESPONSE TIMES

The contribution to 3X74 utilization of a single 3194 display station, when operating at a transaction rate of 400 A-1200 transactions per hour (6.7 per minute), is shown in Figure 34 on page 122 for different environments and modes.

For example, one wants to determine the 3174 utilization and average subsystem response time with thirty-two 3194 workstations in an SNA environment. They have 1920-character screens, and operate in CUT mode at an average rate of 6.7 type A-1200 transactions per minute, with about thirty percent of them active at any one time. From the table in Figure 34 on page 122, we obtain an estimated utilization of 1.21 percent for a single workstation. Therefore,

3174 utilization = 0.30 x (32 x 1.21) = 12 percent

The 3194/3174 subsystem response time for this utilization can now be obtained from the same table by linear interpolation from the response time values in the "10" and "20" percent columns, yielding about 123 milliseconds.

As another example, suppose we have thirty-two 3194s in DFT mode with model 4 screens attached to a 3174 in a non-SNA environment. This time, about eighty percent are active, on the average. The average 3174 utilization would be,

3174 utilization = 0.80 x (32 x 1.46) = 38 percent

and the response time would be 682 milliseconds.

The numbers in the table show that DFT response times rise more slowly with increasing control unit utilization than response times in CUT mode. Although in DFT mode the delay of passing a data stream through is smaller than processing it in the control unit, the response of the 3194 workstation in DFT mode is so large to begin with, that even at high 3174 utilizations, response in DFT mode never drops below CUT mode response.

CU Type	Env Mode	Benchmrk	CU Util, in %	3194 time, msec.	3194/3X74 Response[1], in mseconds, at listed CU utilizations, in %							
					00	10	20	30	40	50	60	65
3174	SNA/CUT	A-1200	1.09	-	98	109	123	140	163	196	245	280
	nSN/CUT	A-1200	1.09	-	98	109	123	140	163	196	245	280
		C-1200	1.31	-	117	130	146	167	195	234	293	334
3174	SNA/DFT	A-1200	0.62	241	296	302	310	320	333	351	379	398
		C-1200	0.49	247	291	296	302	310	320	335	357	373
		A-2160[2]	1.11	434	533	544	558	575	599	632	682	717
3174	nSN/DFT	A-1200	0.82	261	334	342	352	365	383	407	444	470
		C-1200	0.77	267	336	344	353	366	382	405	440	464
		A-2160[2]	1.46	470	601	616	634	657	688	732	798	844
3274	SNA/CUT	A-1200	1.84	-	165	183	206	236	275	330	413	471
		C-1200	2.85	-	255	283	319	364	425	510	638	729
	nSN/CUT	A-1200	1.56	-	140	156	175	200	233	280	350	400
3274	SNA/DFT	A-1200	0.88	241	320	329	340	354	373	389	439	467
		C-1200	0.60	247	301	307	315	324	337	355	382	401
	nSN/DFT	A-1200	0.93	261	344	353	365	380	399	427	469	498

NOTES:

- CU Type — Control unit type: 3174-1L or 3274-41
- Env — Environment: SNA or non-SNA (nSN)
- Mode — CUT or DFT
- CU Util — Control unit utilization, based on 6.7 trans/minute

[1] Data does not include contingency; response times on basis of DE return, add about 20 ms to obtain response to last character

[2] For Model 4 screen size

Figure 34. IBM 3194/3X74 Subsystem Response Times and Utilization

FILE TRANSFER PERFORMANCE (LOCAL CONTROL UNITS)

Refer to Chapter 8 for an introduction to file transfer operation and the use of this performance data.

The equation in Figure 35 is used to approximate the time it takes to receive or send a file. The expansion factor E is larger than 1 when a 3X74 control unit handles also MFI traffic and/or several files concurrently. Chapter 8 shows how the Bc/B ratio is used for estimating E, and the increase in subsystem response time of MFI transactions as a result of one or more concurrent file transfer operations.

For an example of how to use this data, refer to the end of Chapter 8.

Estimated File Transfer Time for 3194 Display Station
A + (B x E) x F seconds

CU	WS	FDV	Size/Mode	Dir	RU/Blk	A	B	FTR	Bc/B
3174 SNA	3194	Dsk	2048/DFT	Rec	1024	4.3	0.22	4.5	0.13
				Snd	256	6.1	0.35	2.9	0.17
3174 nSN			2048/DFT	Rec	2048	2.3	0.32	3.1	0.09
				Snd	2048	2.9	0.39	2.6	0.09
3174 nSN			1920/CUT	Rec		15	0.44	2.3	0.10
				Snd		15	0.45	2.2	0.10

NOTES: # A Initiation + Termination time (seconds)
B Seconds per kbyte of file transferred
FTR File transfer rate (kbytes per second);
for more explanations, see Figure 14 in chapter 8.

Figure 35. 3194 Version 1 - File Transfer Performance Data

PERFORMANCE DATA SHEET FOR IBM PERSONAL COMPUTER WORKSTATIONS

IBM Personal Computer workstations and the IBM Personal System/2[1] Model 30 equipped with an **IBM 3278/79 Emulation Adapter (#5050)** connect to an IBM 3174, 3274, or 3276 control unit. The adapter hardware permits a workstation to be supported in CUT mode or DFT mode. (A 3274 must be customized for DFT mode; the 3276 does not support DFT mode.)

These workstation configurations may use one of several program products. Operating mode (CUT or DFT), and the performance of MFI emulation and file transfer depends on the program support in the workstation and, for file transfer, on the matching support in the host as well. For file transfer operations especially, performance data for one program product should not be used for another. (See Chapter 8 for an introduction to file transfer operations.)

This data sheet considers the performance of the following program products (when attached to local subsystem control units):

- **The IBM 3270 Workstation Program Version 1.0,** enabling the IBM Personal System/2, IBM 3270 PC, and selected models of IBM Personal Computers to emulate CUT and DFT mode display stations. (It replaces and enhances IBM 3270 PC Control Program Version 3.0, among others.) With the 3270PC Host File Transfer Program installed in the host, files can be transferred between the host and the workstation.

- **The IBM PC 3270 Emulation Program, Entry Level,** enabling an IBM Personal Computer (PC, XT, and AT) to emulate 3270 CUT display stations, and file transfer with the 3270PC Host File Transfer Program installed in the host.

- **The Mainframe Communication Assistant[2]** , for 3278/79 emulation (CUT mode only), and file transfer between a PC and a VM/CMS or MVS/TSO host via a 3X74 control unit, or an asynchronous TP connection.

There is an earlier version of the 3278/79 emulation adapter (#2507) which may exhibit slightly different performance characteristics for certain operations. In CUT mode emulation, the #2507 compares to the #5050 as the 3278 to the 3191 display station (both the #5050 and the 3191 accept the new "Read Multiple 32" DCA command).

---

[1] Trademark of the International Business Machines Corporation.

[2] **Trademark of the International Business Machines Corporation.**

IBM 3270 WORKSTATION PROGRAM

The IBM 3270 Workstation Program, Version 1.0 (74X9921), incorporates and enhances the function performed by the IBM PC 3270 Emulation Program Version 3.0. (As discussed in Chapter 7 of the sixth edition of this publication.)

This section provides performance data, for both CUT and DFT modes, on MFI and file transfer operations using IBM 3270 Personal Computer XT and AT, and IBM Personal System/2 Model 30 workstations connected to IBM 3174 Model 1L and 3274 Model 41 control units. For details on workstation hardware, see the applicable IBM 3270 Personal Computer XT and AT and IBM Personal System/2 Model 30 sales pages and publications.

In a CUT mode session, a single 3278 or 3279 model 2 display station can be emulated. In DFT mode, a 3270PC workstations can conduct up to four concurrent sessions with different host programs, provided its 3X74 port has been customized with the appropriate number of device addresses.

Files may be transferred between an application program in the host and a workstation (with the personal computer in DOS command mode), with one of the 3270PC File Transfer Programs below, in addition to VSE/SP 2.1.1 or 2.1.2, or SSX/VSE 1.4.1, installed in the host:

for CMS: # 5664-281 (VM/SP)
for TSO: # 5665-311 (MVS/TSO).

The data is based on using configuration support A, Release 1 for the 3174 Model 1L, and configuration support D, Release 64.1, or later, in the 3274, with RPQ #8K1311, "DFT Display Perform Enhance", for models A. (Without this RPQ, DFT performance will be worse depending on VTAM buffer segment size, more so for smaller sizes, for example, 128 bytes, than for larger ones, for example, 256 bytes).

In DFT mode, these host file transfer programs are capable of transferring files in 2 or 3.5 Kb blocks (by updating them with IND$FIL). The 3.5 Kb blocksize is supported by the IBM 3270 Workstation Program, Version 1, for IBM Personal System/2, IBM 3270 PC, and selected models of IBM Personal Computers. The performance data of the 2 Kb block transfers is based on the installation of the 3270PC Control Program Version 2.1 and PC DOS 3.1 in the workstation.

Customization of the workstation control program determines whether it interacts with the control unit in CUT or DFT mode. Use DFT mode rather than CUT mode, if you have a choice, even when multiple host sessions are not required. For a given transaction load, DFT mode will yield lower control unit utilization than CUT mode. (As noted, only DFT mode permits file transfer in 3.5 Kb blocks, benefitting both transfer rates and control unit utilization.)

Version 1.1 of the IBM 3270 Workstation Program adds support for access from the IBM Token-Ring Network to System 370, and for the workstation to function as a gateway and/or network station for either the IBM PC Network, or the IBM Token-Ring Network. The performance of these functions is not addressed here.

3278/79 EMULATION WITH 3270PC WORKSTATIONS

In **CUT mode,** when emulating MFI transactions as a 3278 Model 2 display station, the 3274-41A/3270PC subsystem response times for the A-1200 and C-1200 benchmarks are about 0.1 second slower than with a 3278 or 3279 display (compare chart 3270-1A* with charts 6-1A and 6-2A in Chapter 6; remember that charts with an asterisk (*) include some contingency, and are based on #2507 adapter performance.) This delay is due to the transfer and updating of the screen contents by the PC control program.

Note: The subsystem response times are based on the return of Device End to the host. Up to 0.1 seconds, depending on workstation type, may elapse before the last character is written on the screen.

Because of their faster microprocessor, the response of 3270 Personal Computer AT workstations is about the same as for 3278s. See chart 3270-1B*.

The MFI performance improvement obtained when attaching to a 3174 subsystem is comparable with the performance obtained for a 3278 (3191), as depicted in charts 6-1A through 6-3B (Chapter 6). See also Figure 7 on page 44.

The display monitor selection usually does not affect response time, except the 3295 Plasma Monitor, which increases response time about 0.4 second.

Running PC programs in the workstation background does not affect MFI transactions, because the latter have higher priority. Conversely, execution of PC programs may be slowed somewhat as a result of MFI operations.

A 3X74 control unit interacts with a 3278/79 emulation adapter (#5050) as with a 3191 display station (for #2507, as with a 3278). Therefore, the effect of a workstation emulating a 3278, on the operations of other devices attached to a 3X74 will be about the same as with actual 3278 (3191) display stations.

For **DFT mode,** charts 3270-1A* and -1B* show the corresponding 3274/3270 PC subsystem response curves for the A-1200, C-1200, and C1-1200 benchmarks.

The subsystem response characteristics for the 3270PC and AT in DFT mode resemble those for 3290 attachment. (Compare charts 6-4*, 6-6A*, and 6-6B* in Chapter 6.) Note that the 3270PC exhibits somewhat better response times than the 3290, and that the AT is even better. Also, types C and C1 benchmarks perform better relative to the type A benchmark in DFT mode than in CUT mode.

The effect of DFT operation on the response time of other 3X74-attached terminals is about the same as with 3290 display stations.

The performance of a PC application running in the background will be impacted somewhat more in DFT mode than in CUT mode because in DFT mode the 3270PC must process the data stream.

In a 3274 cluster mixing 3278s and 3270PCs operating in DFT mode, the subsystem response characteristics shown in chart 3270-2A* are comparable with those of a mix of 3278s and 3290s (chart 6-5* in Chapter 6).

Chart 3270-2B* illustrates how a fast workstation processor can show up the DFT mode advantage by providing reasonable responses for large transaction loads.

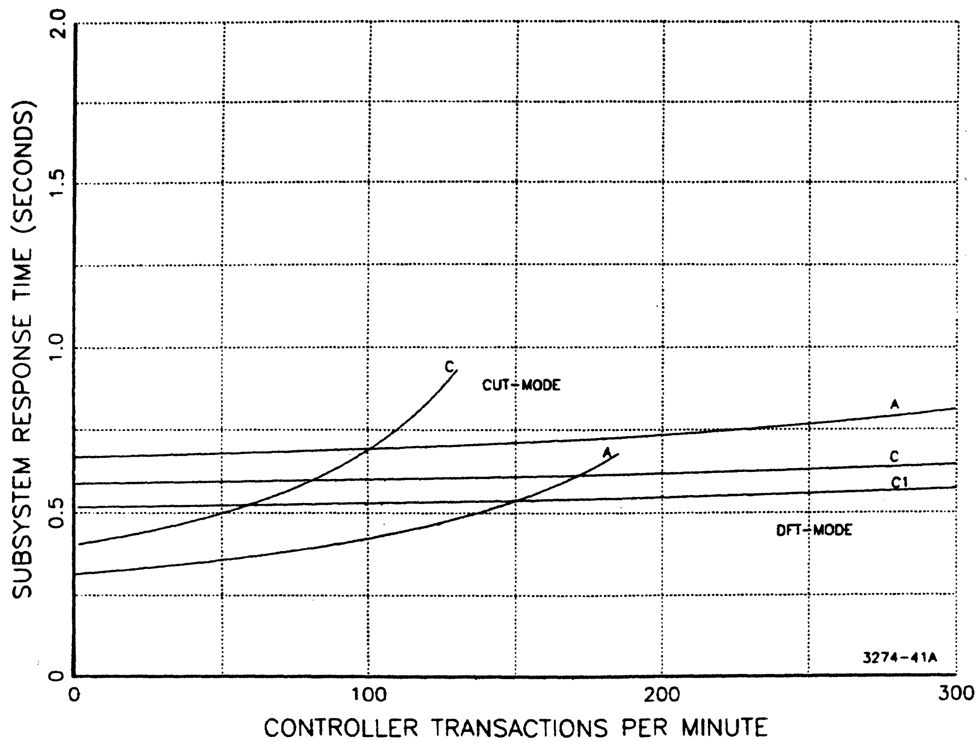

Chart 3270-1A*: CUT/DFT Response of 3270PC on 3274-41, A/C/C1

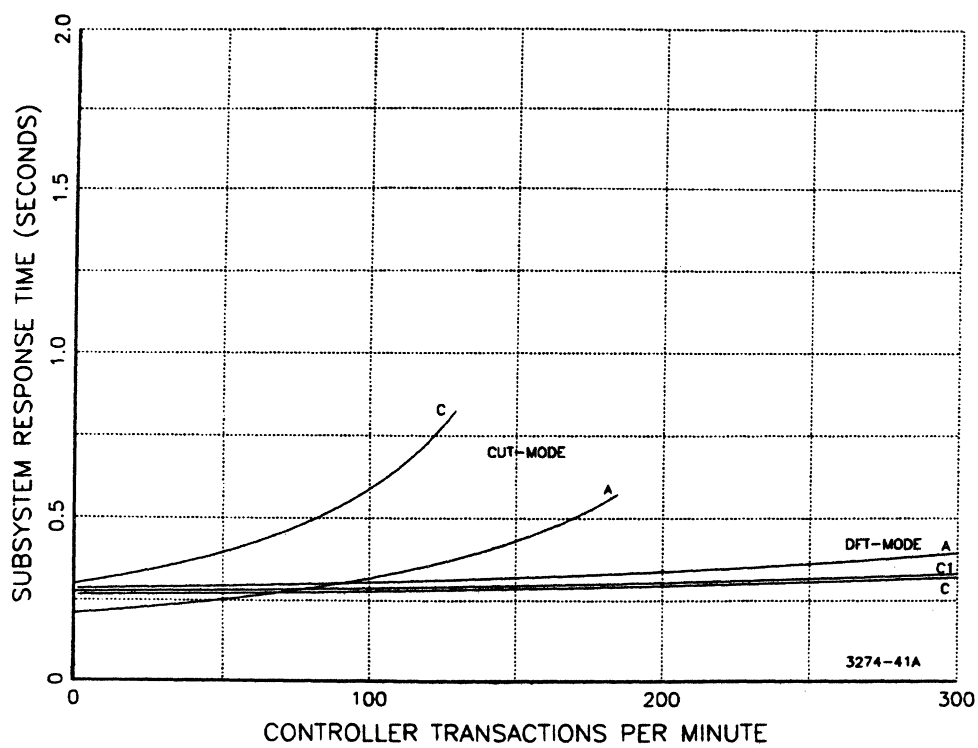

Chart 3270-1B*: CUT/DFT Response of 3270-AT on 3274-41, A/C/C1

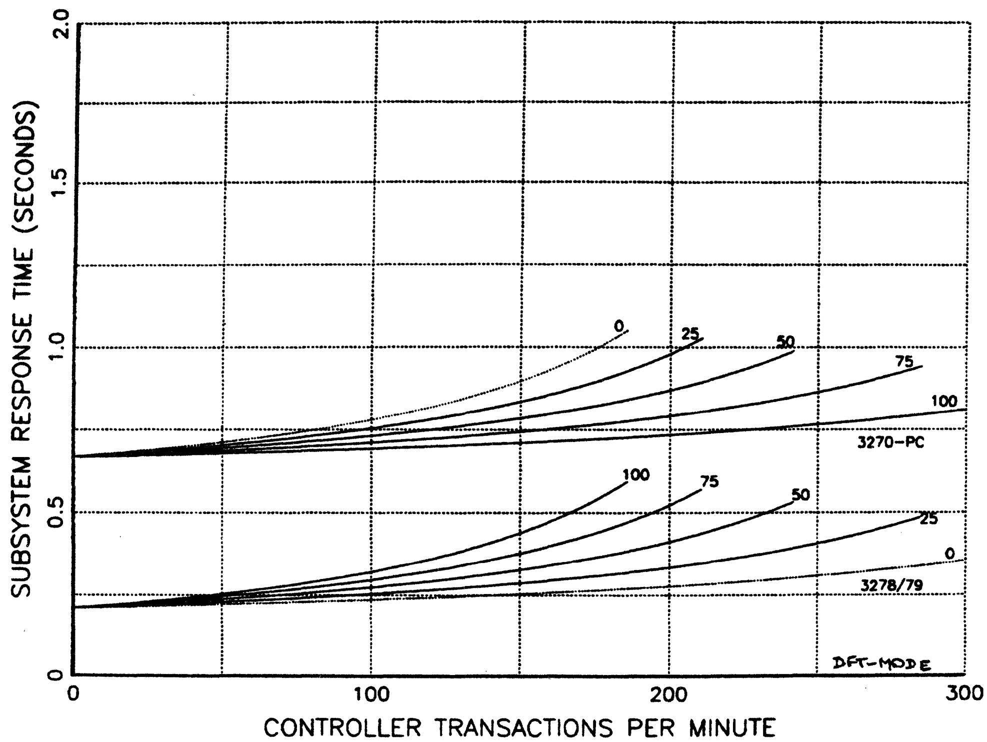

Chart 3270-2A*: Response of 3270PC(DFT)/3278 Mix on 3274-41

3274-41 WITH 3278/79/3270PC-AT MIX A DATA STREAM

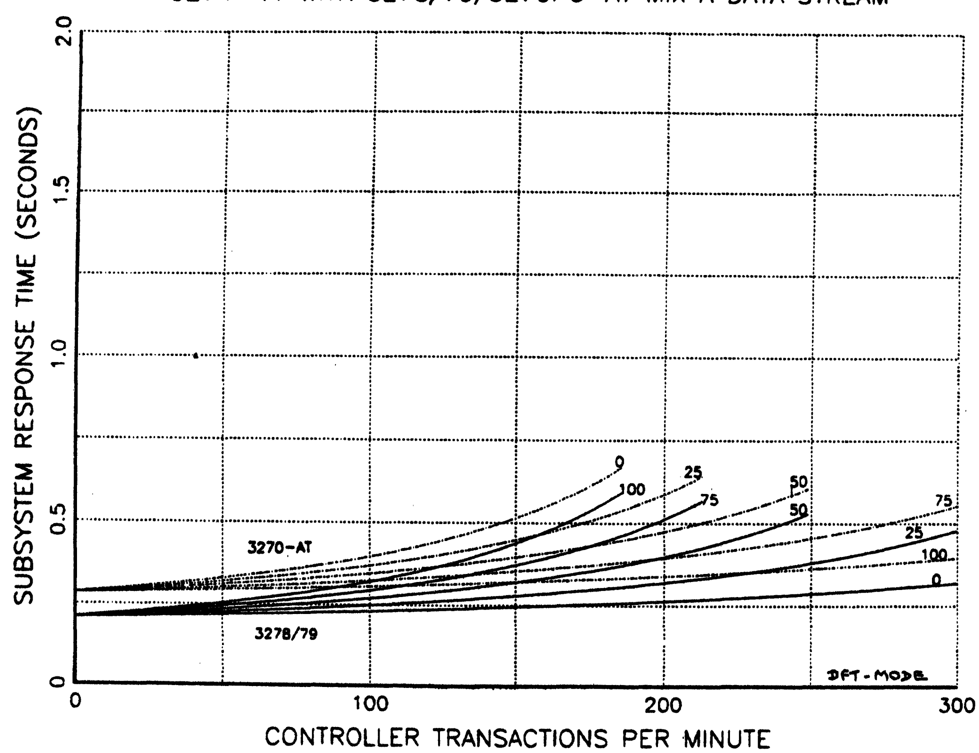

Chart 3270-2B*: Response of 3270-AT(DFT)/3278 Mix on 3274-41

FILE TRANSFER PERFORMANCE (LOCAL CONTROL UNITS)

Outbound data is transferred from the host to the workstation by the Receive command (using the ASCII option). Data is transferred inbound by the Send command (using the ASCII and Carriage Return Line Feed options (CRLF)). Be sure that in CUT mode the File Transfer Aid bit in configuration question 125 for the 3174 control unit has been set to "1". The monitor type used with the workstation does not affect file transfer performance.

The ASCII command option causes the translation between ASCII and EBCDIC to be performed by the file transfer program in the host. The CRLF option recognizes carriage return/line feed characters, and transforms them in appropriate record separators before storing the file in host storage.

File size depends on the application and varies greatly. In CUT mode, in both SNA and non-SNA, a file is transferred, in either direction, in blocks of about 1920 bytes; in DFT mode, in blocks of 2 Kb or 3.5 Kb.

Remember to set an appropriate value in the CONFIG.SYS file, for example, BUFFERS=10 for starters, to ensure a large enough RAM area for accommodating an adequate number of these blocks, which avoids the need for frequent access to the file device and therefore helps maximize performance.

Figure 36 on page 132 lists, for various cases, the constants and coefficients in the equation for estimating minimum file transfer time as a function of file size. The "Bc/B" ratio column specifies the appropriate entry in Figure 12 on page 61 for obtaining the expansion factor E which accounts for the number of concurrent file transfers and/or MFI transaction load being processed by the control unit.

File transfer performance depends on many parameters, such as operational mode (CUT or DFT), whether it is a Receive or Send operation, the host response times, and whether a diskette drive or fixed disk in the workstation is involved. The performance data assumes that the file device has adequate storage space that is not highly fragmented. File fragmentation will reduce file transfer performance.

Host response time depends on mainframe speed, host processing load, and the configuration of the operating system. For example, running MVS under VM may double the host response time obtained with MVS running by itself. For the conditions reflected by the numbers in the table, this would result in an increase of up to 15 percent in file transfer time.

For estimating the increase in MFI response times as a result one or more concurrent file transfers, refer to the example in Chapter 8.

File Transfer Performance of Personal Computer Workstations
Estimated time: A + (B x E) x F seconds

CU	WS	FDV	Size/Mode	Dir	RU/Blk	A	B	FTR	Bc/B
3174 nSN	XT	RAM	1920/CUT	Rec		9	0.68	1.5	0.07
				Snd		9	0.70	1.4	0.07
			3584/DFT	Rec		10	0.22	4.6	0.11
				Snd		10	0.24	4.1	0.10
SNA			3584/DFT	Rec		12	0.19	5.2	0.13
				Snd		13	0.28	3.5	0.09
3174 nSN	AT 339	RAM	1920/CUT	Rec		6	0.17	5.8	0.27
				Snd		6	0.18	5.5	0.25
			3584/DFT	Rec		4	0.13	7.6	0.19
				Snd		5	0.12	8.2	0.21
SNA			3584/DFT	Rec	1024	7	0.13	7.5	0.19
				Snd	256	8	0.20	4.9	0.13
3174 nSN	PS/2 Mod 30	RAM	1920/CUT	Rec		10	0.35	2.9	0.13
				Snd		11	0.39	2.6	0.12
			3584/ FT	Rec		9	0.13	7.5	0.19
				Snd		10	0.15	6.7	0.17
SNA			3584/DFT	Rec		11	0.13	7.5	0.19
				Snd		11	0.23	4.4	0.11
3274-41	PC R1.2	Fdk[1]	1920/CUT	Rec		14	.75 (.88)	1.3	0.15
				Snd		14	.75 (.88)	1.3	0.25
			2048/DFT	Rec		10	.40 (.43)	2.5	0.19
				Snd		11	.25 (.29)	4.0	0 43
3274-41D	AT R2.1	Fdk[1]	1920/CUT	Rec		2	.34 (.53)	2.9	0.39
				Snd		2	.36 (.41)	2.8	0.49
			2048/DFT	Rec		2	.14 (.27)	7.1	0.56
				Snd		2	.14 (.28)	7.1	0.33
3274-41A	AT R2.1	Fdk[1]	2048/DFT	Rec		3	.22 (.34)	4.6	0.56
				Snd		4	.29 (.38)	3.5	0.33
3274-41D	PC Version 3.0	Fdk	1920/CUT	Rec		3	0.10	10	0.60
				Snd		3	0.10	10	0.62

NOTES:
- A — Initiation + Termination time (seconds)
  B — Seconds per kbyte of file transferred
  FTR — File transfer rate (kbytes per second), for more explanations, see Figure 14 in Chapter 8.

[1] B-values in parentheses for workstations with diskette

Figure 36. File Transfer with IBM 3270 PCs, and PS/2 Model 30

THE IBM PC 3270 EMULATION PROGRAM, ENTRY LEVEL (3274 DATA ONLY)

An IBM Personal Computer (PC, XT, AT) equipped with an IBM Personal Computer 3278/79 Emulation Adapter (#5050 or #2507), and using the IBM PC 3270 Emulation Program, Entry Level (59X9904) running under DOS 2.1, or higher (PC and XT), or DOS 3.1, or higher (PC, XT, AT), can emulate many of the functions of an IBM 3278 Model 2 or, with an IBM 5272 Color Display installed, a 3279 Model 2A or S2A. (For more information and limitations, see the sales pages and the IBM PC 3270 Emulation Program, Entry Level User's Guide.)

The transfer of files in either direction between an application program in the host, and main memory, diskette, or fixed disk in the workstation requires, in addition, the 3270PC Host File Transfer Program (5664-281 for VM/SP or 5665-311 for MVS/TSO) and VSE/SP 2.1.1 or 2.1.2, or SSX/VSE 1.4.1 running in the host.

The effect of running a PC DOS application in the PC background on the performance of transactions with the host (3278/79 emulation and file transfers) is negligible, because the host transactions have a higher priority. Thus, the PC DOS application may encounter delays depending on the nature and frequency of host interactions.

FILE TRANSFER OPERATION (LOCAL CONTROL UNITS)

Refer to Chapter 8 for an introduction to file transfer operation, and the example at the end for the use of this performance data.

The data transfer rate between the host and a workstation by a Receive (download) or Send (upload) command can be estimated with the equation and data in Figure 37 on page 134. (Data does not apply when the 3270 emulator is running under TopView[3] .) The coefficient B is used to estimate the minimum time it takes to receive or send a single file with no other transactions in the control unit. (The expansion factor E =1.)

The factor E is larger than 1 when a 3X74 control unit also handles MFI traffic and/or several files concurrently. Chapter 8 shows how the Bc/B ratio is used for estimating E (Figure 12 on page 61), and the increase in subsystem response time of MFI transactions as a result of one or more concurrent file transfer operations.

A slow or heavily loaded host can slow down file transfers, while transfer times may also be extended by a highly fragmented, near-fully loaded file device in the workstation.

The significant differences between Send and Receive can be explained by the fact that for the transfer of the same amount of data, sending requires more blocks than receiving.

---

[3] Trademark of the International Business Machines Corporation.

File Transfer with PC 3270 Emulation Program, Entry Level, Version 1.0
Estimated time: A + (B x E) x F seconds

CU	WS	FDV	Size/Mode	Dir	RU/Blk	A	B	FTR	Bc/B
3174 nSN	AT	Fdk	1920/CUT	Rec		5	0.17	5.9	0.60
				Snd		8	0.50	2.0	0.60
3274-41D	AT	Fdk	1920/CUT	Rec		6	0.25	4.0	0.60
				Snd		10	0.55	1.8	0.60
3174 nSN	XT	Fdk	1920/CUT	Rec		6	0.53	1.9	0.30
				Snd		9	0.91	1.1	0.30
3274-41D	XT	Fdk	1920/CUT	Rec		6	0.53	1.9	0.30
				Snd		9	0.84	1.2	0.30

NOTES:
- A Initiation + Termination time (seconds)
  B Seconds per kbyte of file transferred
  FTR File transfer rate (kbytes per second),
  for more explanations, see Figure 14 in Chapter 8.
- Data for ASCII text files; transfer in blocks < 2 kbytes
- WS diskette lightly loaded, not higly fragmented.

Figure 37. File Transfer Performance of IBM 3270 PC Emulator, Entry Level

IBM MAINFRAME COMMUNICATION ASSISTANT PROGRAM (3274 DATA ONLY)

The IBM Mainframe Communication Assistant (6024140) consists of two programs, one for installation on an IBM Personal Computer or IBM Personal Computer XT, and another for installation on an IBM mainframe. The PC part of the IBM Mainframe Communication Assistant runs under DOS 2.0 or higher, and can run under TopView with full windowing capability.

The PC or PC XT must be equipped with the 3278/79 emulation adapter for connection to a 3X74 control unit, and an Asynchronous Communication Card and a modem (or equivalent) for the asynchronous line.

IBM Mainframe Communication Assistant is a member of the IBM Assistant Series, and allows integration of mainframe data into IBM Filing Assistant files. That same information can then be ported to the other members of the IBM Assistant Series family of PC software products.

This program product supports direct connection to a VM/CMS or MVS/TSO host via a 3X74 control unit, and connection via an asynchronous TP link to a host and/or other PCs running the IBM Mainframe Communication Assistant software. With the 3X74 connection (CUT mode only), the program permits menu selection of 3278/79 model 2 emulation (U.S. English keyboard, 4-color only), and the up- and downloading of files. With an asynchronous connection, the program allows IBM 3101 emulation and transferring files over this link as well.

This section addresses the performance of 3278/79 emulation and file transfer through a local 3274 control unit only. The Mainframe Communication Assistant can transfer changes made to files between your IBM PC and a host computer. A file may be edited on your PC and then just the added or changed lines will be sent to the host computer.

The program is provided on three dual-sided disks: a Mainframe Communication Assistant PC Programs diskette, a Mainframe Communication Assistant Host Programs diskette, and a Mainframe Communication Assistant Sample Programs diskette.

For more information on the operations of this program product, refer to the user's guide for the IBM Mainframe Communication Assistant program.

UP- AND DOWNLOADING OF FILES (LOCAL CONTROL UNITS)

When transferring a file, you must specify whether it is an ASCII text file or a binary (or XASCII: Extended ASCII) file. In addition to uploading and downloading of files, the program allows changes in ASCII text files to be uploaded, but not downloaded.

To transfer a file, the host/PC file transfer programs divide a file into a series of data blocks of approximate screen buffer size.

The information in Figure 38 may be used to approximate the minimum time it takes to receive (download) or send (upload) a file. They assume a lightly loaded 3081 host.

File Transfer with IBM Mainframe Communication Assistant Estimated time: A + (B x E) x F seconds										
CU	WS	FDV	Size/Mode	Dir		A	B	FTR	Bc/B
3274-41D	PC	Dsk	1920 CUT	Rec		18	0.9	1.11	0.10
				Snd		19	1.2	0.83	0.10

NOTES:

- A Initiation + Termination time (seconds)
  B Seconds per kbyte of file transferred
  FTR File transfer rate (kbytes per second), for more explanations, see Figure 14 in Chapter 8.
- Data for ASCII text files; transfer in blocks < 2 kbytes
- WS diskette lightly loaded, not higly fragmented.

Figure 38. File Transfer Performance of IBM MC Assistant Program

The utilization of a 3X74 control unit by this file transfer protocol is relatively small. A file transfer may take longer than the estimated time with a slow or heavily loaded host, and a rather full, highly fragmented diskette in the workstation.

Refer to Figure 12 on page 61 and associated text for determining the expansion factor E from the applicable Bc/B ratio in Figure 38 and the average control utilization U as determined by the MFI traffic generated by other active workstations. (For a single file transfer with CTR=0, E=1 by definition.)

The data in the table reflect the observations for concurrent uploading of files. For download operations, the E-values may be slightly less, but not more than five percent of the amount in excess of 1.

To determine the 3X74 subsystem response time increases as a result of one or more concurrent file transfers, use the information in Figure 13 on page 62 and associated text with the example.

PERFORMANCE DATA SHEET FOR IBM 3X74 WITH 4245 LINE PRINTERS

The IBM 4245 Model D20 and D12 line printers attach to IBM 3174 and 3274 subsystem control units. (They also attach to other systems, but only 3X74 attachment data is included here.)

The information in this performance data sheet should be used in conjunction with Chapter 11, which introduces the various printer performance characteristics.

Supported data streams are LU types 1 and 3 in SNA, and DSC in non-SNA environments. The EPC function may be invoked for DSC and LU type 3 operations. These printers support decompression and decompaction in LU type 1.

The maximum print speed (LPM) for the D12 and D20 models depends on the type of print band used in the printer. (See Figure 39 on page 138.) Throughput is slightly affected by the number of lines per inch (LPI) being printed, and the percentage of blank lines being skipped. Printing with 8 LPI rather than 6 LPI may increase this rate in the order of 3 to 8 percent, while the presence of blank lines will decrease throughput somewhat.

To achieve maximum performance, it is recommended to use the NL control character for moving to a new line. Avoid using the character sequence CR,NL in a data stream, because it may reduce performance appreciably.

MAXIMIZING THROUGHPUT

In SNA environments, use LU type 1 data streams for best performance of a 3X74/4245 printer subsystem. The recommended RU size is 1024 bytes with pacing =2, and exception response. In a specific application, it may be possible to improve on this by some fine tuning.

For remote 3X74 subsystems, use compression or compaction, when supported in the host.

In non-SNA environments, only 3270 data stream compatibility mode (DSC) can be used. Print with the EPC function turned on, unless it presents unacceptable recovery problems. Note that the 3174 control unit with Configuration Support-A Release 3 microcode supports the "Set Printer Characteristics Structured Field" function, enabling a host program to set and reset the EPC parameter. In DSC (and LU type 3) mode, especially at high print rates, printer throughput reduction will be severe with EPC turned off.

The use of printer adapter buffer sizes larger than 1920 bytes (model 2 emulation) is not likely to improve performance of 4245 printers. Unloading time will be extended because the buffer space in the printer itself is also in the order of 2 Kb.

For LU type 3, being the approximate functional equivalent of DSC in SNA environments, follow the recommendations for DSC.

PERFORMANCE INFORMATION

Figure 39 on page 138 lists average data rates (in kbytes per second) for the various 4245 model/print band combinations, at three line widths. For example, in cases where we know that the actual throughput is less than the model/print band combination indicates, such as for DSC and LU type 3 with the EPC turned off, use the data on the line with a matching LPM. (See section on "DSC and LU type 3 operation" in Chapter 11 for determining throughput decrease.)

In addition, average 3174 and 3274 control unit utilization percentages for LU type 1 and DSC mode have been computed (based on measurements, no contingency included).

High utilization of a 3X74 control unit will not cause it to fail, but may occasionally introduce delays large enough to halt printing, thereby reducing throughput.

		Average Line Width								
		130-character			100-character			70-character		
	LPM	Rate kby/s	3174 U %	3274 U %	Rate kby/s	3174 U %	3274 U %	Rate kby/s	3174 U %	3274 U %
FOR IBM 4245 D20:										
48 char. band	2000	4.40	11/12	32/23	3.40	09/09	25/18	2.40	06/07	18/13
50/54/63/64	1570	3.45	09/09	25/18	2.67	07/07	20/14	1.88	05/05	14/10
94/98/116	1051	2.31	06/06	17/12	1.79	05/05	13/10	1.26	03/03	10/07
108/124/127/142	744	1.64	04/04	12/09	1.26	03/03	09/07	0.89	02/02	07/05
FOR IBM 4245 D12:										
48 char. band	1270	2.79	07/07	20/15	2.16	06/06	16/12	1.52	04/04	11/08
50/54/63/64	980	2.16	05/06	16/11	1.67	04/05	12/09	1.18	03/03	09/06
94/98/116	640	1.41	04/04	10/07	1.09	03/03	08/06	0.77	02/02	06/04
108/124/127/142	445	0.98	02/03	07/05	0.76	02/02	06/04	0.53	01/01	04/03

NOTES:

- Print speeds at 6 lines per inch (LPI)
- Line widths exclude control characters
- kby/s: Data rate in kbytes/second
- U % : Approximate 3X74 utilization with LU type 1 and DSC

Figure 39. 3X74 Utilizations For 4245 Printer Operations

Charts 4245-1A and -2A show how average 4245 throughput decreases with increasing MFI activity for the 3174 and 3274, respectively. The lesser sensitivity of the 3174 is due to its smaller utilization for a given throughput (data rate).

EFFECT OF 4245 OPERATION ON WORKSTATIONS RESPONSE

The effect of a 4245 printer on the response of workstations attached to the same control unit depends on the average data rate to the printer, as shown in charts 4245-1B and -2B.

For another way to assess this effect, note that, for example, one 3278-2 display station processing an average 6.7 type A-1200 transactions per minute utilizes about one percent of a 3174-1L control unit, or 16 percent for 16 active displays.

When a 4245 D20 with 48-character print band printing 132-character lines adds about 12 percent (Figure 39 on page 138) utilization, the total 3174 increases to about 28 percent.

With the data in Figure 7 on page 44 (Chapter 6), we determine (by interpolation) that 3174/3278 average non-SNA subsystem response time increases from about 147 to 167 milliseconds with a 3174 utilization increase from 16 to 28 percent.

RECOMMENDATIONS

In planning a 3X74 subsystem configuration that includes a 4245 printer with high to medium throughput, it is recommended to plan for less than 50 percent utilization to maintain printer throughput near (80-90 percent) maximum.

For remote 3X74 control units, the line speed rather than the control unit limits throughput. Using the throughput numbers in the table again, one may obtain a rough idea about the ability of the telecommunication line to keep up with the traffic requirements generated by the devices attached to the remote control unit. Plan to keep the aggregate utilization of an SDLC or full duplex BSC line below 50 percent, and a half duplex BSC line below 30 percent. (See also Chapter 9.)

With respect to attaching workstations with file transfer capability to the same control unit as a 4245 printer, there are several things to consider:

- The workstation type
- The number of attached workstations
- The size and frequency of files being transferred.

Control unit utilization as a result of file transfer is dependent on workstation type because their file transfer rates differ. Note, for example, that some workstations are capable of file transfer rates of up to 10 kbytes per second in 3.5kbyte blocks. Such file transfer operations can generate appreciable control unit utilizations, and are therefore likely to considerably affect the operation of printers with high data rate requirements. With the continuing improvements of file transfer rates, the trend is for these effects to increase.

The recommendation is that when high-throughput 4245 printers are attached to a 3X74, exercise caution in attaching and using such workstations, preferably limit attachment to those with low transfer speeds, and use their file transfer capability infrequently.

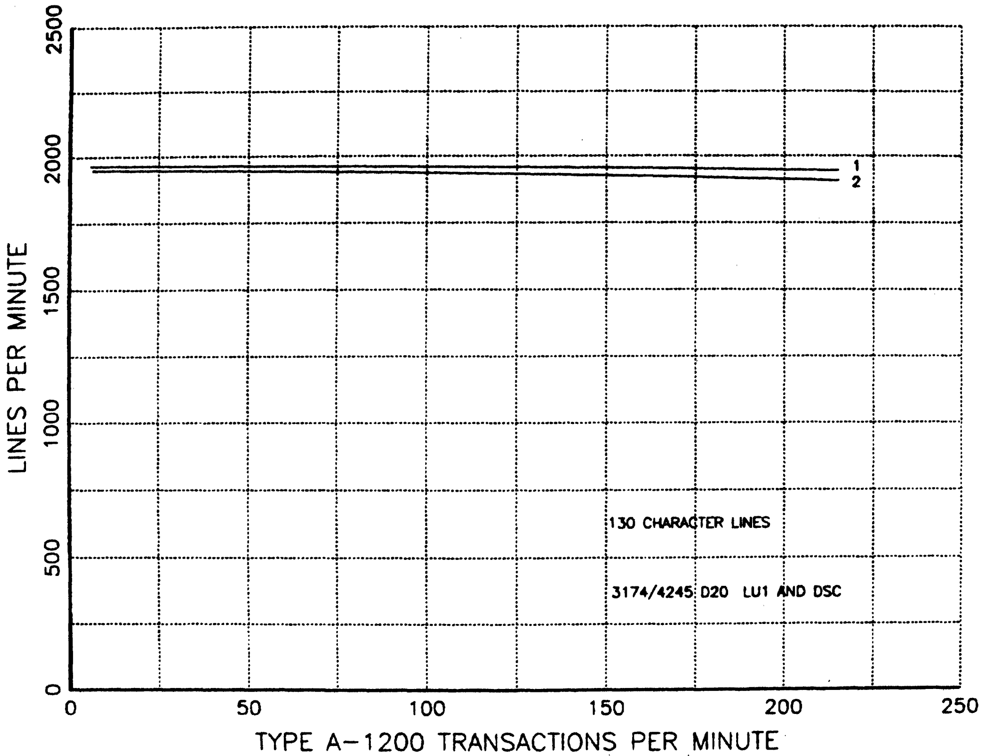

Chart 4245-1A: 4245 Print Time on 3174 with 3278s

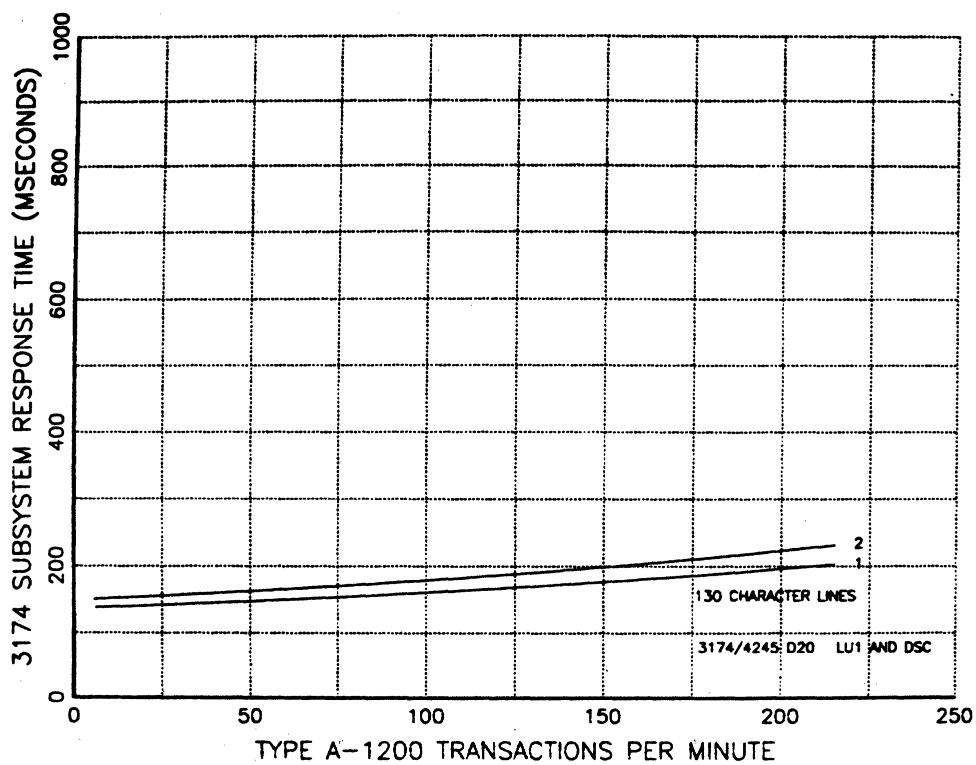

Chart 4245-1B: 4245 Effect on Response of 3174/3278

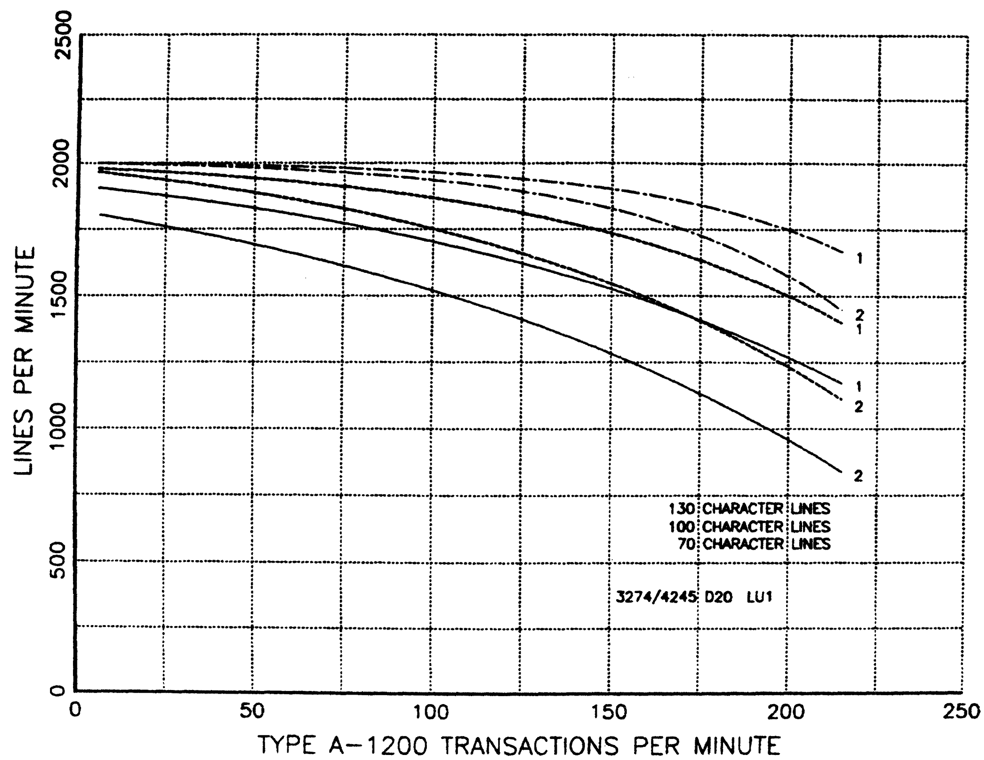

Chart 4245-2A: 4245 Print Time on 3274 with 3278s

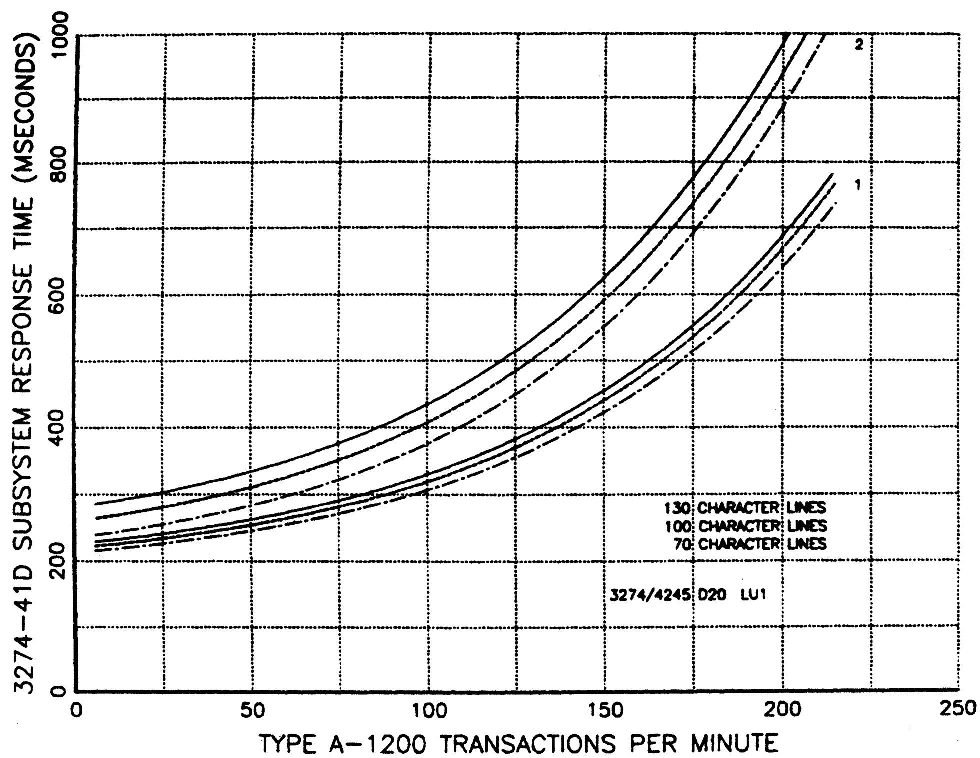

Chart 4245-2B: 4245 Effect on Response of 3274/3278

PERFORMANCE DATA SHEET FOR IBM 4250 DOCUMENT PRINTER

The IBM 4250 document printer uses an electro-erosion technology. A print head with a row of thirty-two electrodes moves from left to right, eroding the coating of its special paper on its way. During the return of the print head to its starting position, the image data for the next sweep is prepared, and then the paper is advanced. The print resolution is 23.6 dots per mm (600 dots/inch).

The 4250 attaches to both SNA and non-SNA 3X74 control units, using their printer protocol with some minor enhancements. (Configuration support C or D is required for the 3274.) Because of the high average data rate required by this printer (on the order of 4 kbytes per second), it is recommended to use a 56 kbps line for remote control units (available with SNA units only), and severely limit the attachment of other devices.

The 4250 print head prints a 32 dot high image swath of specified length (472< length <7020 dots) using decompressed image data residing in the sweep data buffer. The actual sweep length is somewhat larger than obtained by dividing the line length (in dots) by 600, because some travel is required for head acceleration and deceleration. Print head speed is constant at 1000 mm per second (39.3 inches/second).

The outbound printer data stream contains print image data in compressed form. The host provides this data on demand on a per sweep basis. The data passes through the 3X74 control unit and is buffered in an 8 Kb buffer associated with the coax adapter in the 4250 printer. In non-SNA, the host is programmed to send single messages less than 4 Kb. The control unit uses the coax adapter buffer as two 4 Kb buffers in flip-flop fashion. In SNA, the data for a single sweep is transmitted as chained RUs. Each RU is 1536 bytes or less.

After a sweep has been printed and the print head starts to reverse direction, the compressed sweep data for the next sweep in the coax adapter buffer is transferred into the 32 Kb sweep buffer with decompression "on the fly". The paper may be advanced beyond the minimum of 32 dots in 8-dot increments whenever there is no information to be printed in the intervening space. Note that the 4250 printer operates asynchronously, that is, a sweep will not start until loading of the sweep data buffer and paper advance has been completed.

PERFORMANCE OVERVIEW OF THE 4250 PRINTING SUBSYSTEM

The 4250 printer is a high resolution, all points addressable device. This results in large quantities of image data being printed on a single page (possibly as much as 30 megabits), which in turn would require a very large data flow through a 3X74.

To reduce data traffic, the data stream is compressed in the host and decompressed in the 4250. The actual data rate is a function of the document characteristics which determine the achievable degree of compression.

From a performance point of view, there are two aspects worth evaluating:

1. At what point will the combined MFI transaction traffic cause a 4250 printer to operate at less than maximal rate?

2. To what extent will the operation of 4250 printer(s) limit the combined transaction rate of MFI display stations on the control unit, and slow their response?

The performance of the 4250 printer is expressed in terms of time to print a page, as opposed to response time used to specify MFI terminal performance. The 4250 performance is dependent upon:

- Size of the page

- Content of the page.

In addition, its performance may be affected by:

- Type of control unit

- Load on control unit caused by other printers and displays

- Availability of host to provide data.

The size of the page and of the area(s) where information is printed determines performance parameters such as number and length of sweeps. The frequency and length of skips depend on the amount of area on a page left blank as a result of spacing.

As noted before, the amount of data transmitted for a sweep depends on sweep-length as well as on the achievable compressibility ratio.

For the purpose of this guide, a typical A-sized page with the following characteristics was selected:

```
Page length....................... 280 mm (11 inch)
Page width........................ 216 mm (8.5 inch)
Number of sweeps (across width).... 135
Average sweep length L........... 3726 dots
Average sweep data size.......... 2586 bytes
Data content, per page......... 349,110 bytes
Average compression factor........ 5.76
Skip frequency.................. 33.8 percent
Average skip length................ 66 dots
```

The 4250 printer will print this page in 85 seconds, that is, at an average rate of 95 sweeps per minute. The average data rate through the control unit at this print speed is about 4.1 kbytes per second.

Depending on the document being printed, these numbers will be higher or lower depending on size and content. For example, the total time for executing a sweep is not constant. It can be estimated, in milliseconds, by using the expression:

$$Tsweep = 300 + (50.8 \times L/600)$$

This expression is only valid if the distance over which the paper is advanced (skip length) is reasonably small.

The amount of compressed data to be passed through the 3X74 control unit to print a sweep depends on its image content, and may vary from less than a hundred to as many as 28000 bytes for incompressible data. For most documents, however, anywhere from a few hundred bytes to 6000 bytes per sweep appears to be a realistic expectation.

With respect to host availability, the assumption is that adequate resources (host capacity and main memory) have been allocated to operate a 4250 printer at its maximum speed, that is, to respond to its demand for data with minimal delays. In this context, it should be noted that the Composed Document Printing Facility Program Product used for driving the printer maintains data for sixteen sweeps in main storage to avoid delays in the host.

PERFORMANCE CRITERIA

For a 3X74-based subsystem with a 4250 document printer attached, the primary performance criterion is that the 4250 can perform at maximum speed (as determined by document size and content), irrespective of operations on other displays or printers attached to the same control unit.

A secondary criterion is that the operation of one or more 4250 printers at maximum rate will not slow the response of displays on the same control unit to unacceptable levels.

Chart 4250-1A shows how the print time of the benchmark page is affected by attachment of one or more 4250 printers and transaction rate on displays attached to the same 3174-1L or 3274-41A control unit.

Chart 4250-1B depicts how subsystem response time curves for displays are affected by one or more 4250 printers attached to the same 3174-1L or 3274-41 control unit respectively.

The equivalent non-SNA operations exhibit about the same performance characteristics.

The curves clearly exhibit the improved 3174 control unit performance by the lesser effect of printers and displays on each other in comparison with the 3274.

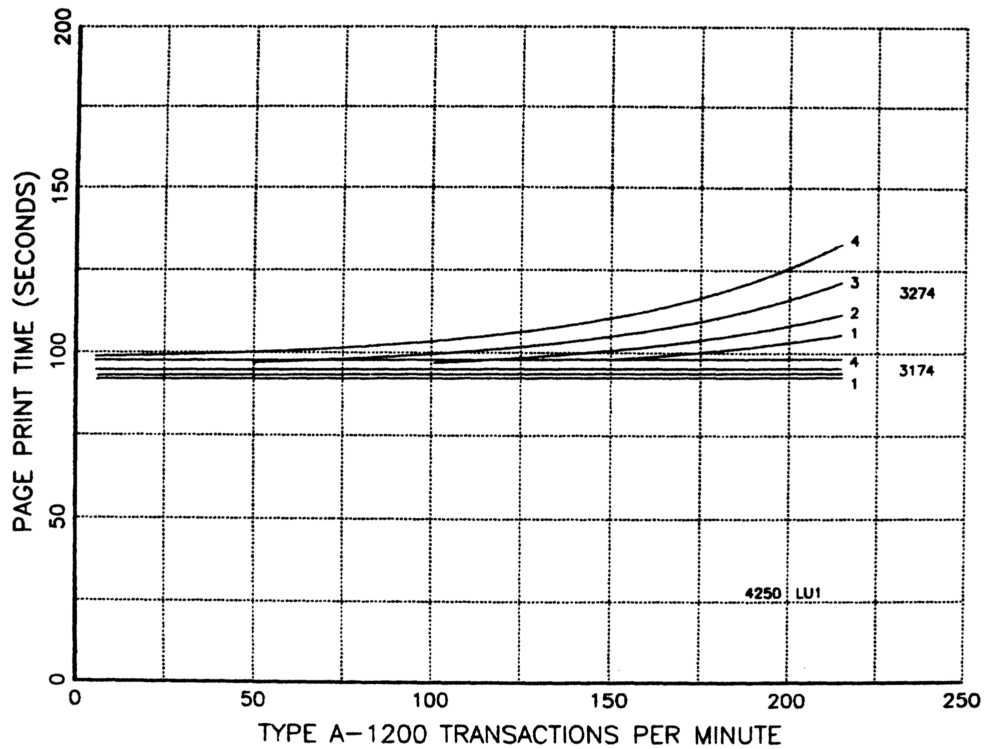

Chart 4250-1A: 4250 Page Print Time on 3X74 with 3278s

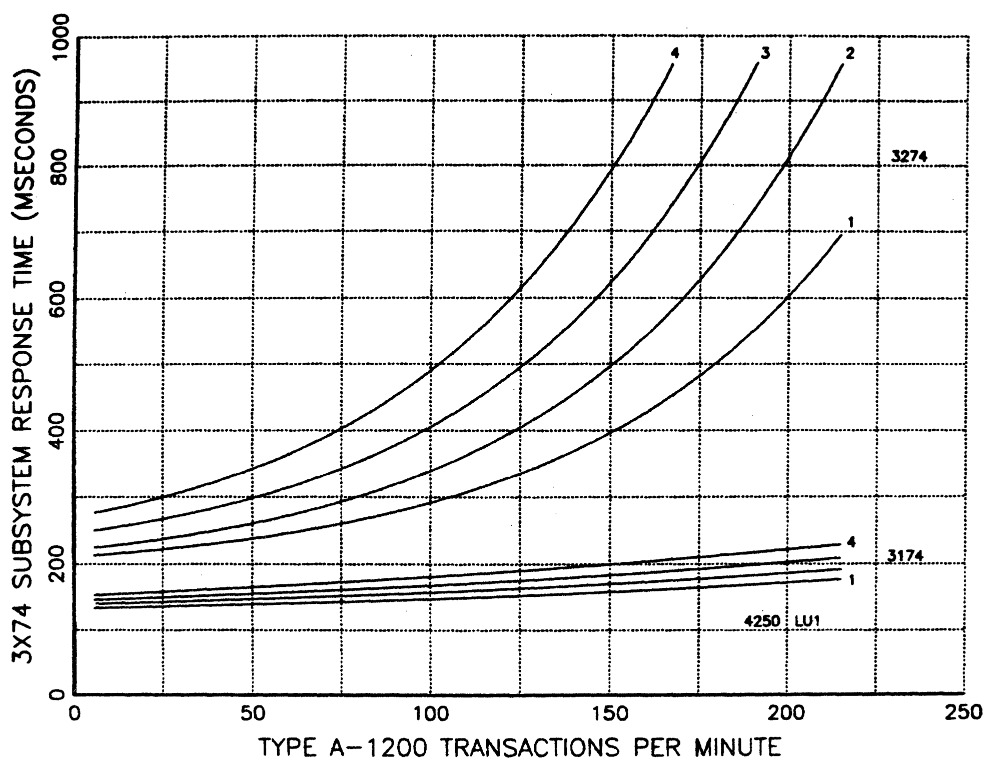

Chart 4250-1B: 4250 Effect on Response of 3X74/3278

PERFORMANCE DATA SHEET FOR RT PERSONAL COMPUTER (3274 DATA ONLY)

An IBM RT Personal Computer workstation (based on the IBM 6150 or 6151 System Unit) equipped with an IBM Personal Computer Emulation Adapter (#5050 or #2507) can be connected to an IBM 3174 or 3274 subsystem control unit. With the IBM RT Personal Computer AIX operating system installed, the IBM RT Personal Computer 3278/79 Emulation program (5669-052) can emulate, with some limitations, 3278-2, 3279-2A, or -S2A display stations with a 3278 keyboard (US English).

3278/79 Emulation supports the transfer of files between the IBM RT Personal Computer and an application program in a VM/CMS or MVS/TSO host. The 3270-PC Host File Transfer Program (5664-281 for VM/SP or 5665-311 for MVS/TSO) and VSE/SP 2.1.1 or 2.1.2, or SSX/VSE 1.4.1 needs to be operative in the host.

The effect of running a concurrent RT Personal Computer or Coprocessor application on the performance of MFI transactions and file transfers is negligible, because host transactions have a higher priority. In addition, an RT PC application can be run at the same time.

For example, when interacting with the RT PC application, the user can switch to 3278/79 emulation, initiate a host transaction, and switch back to the RT PC application. The host will continue execution of the transaction and return the response, if any, to the personal computer where it is stored and available for viewing. Thus, both sessions are processed concurrently, although only one can use the display screen at a given time.

For more detail on these operations, refer to **IBM RT Personal Computer 3278/79 Emulation,** SV21-8032.

The 3278/79 emulation performance of the RT Personal Computer is the same as, or better than that of a 3270 Personal Computer AT (CUT mode). (See chart 3270-1B* in the performance data sheet for 3270 PC workstations.)

FILE TRANSFER PERFORMANCE (LOCAL CONTROL UNITS)

Refer to Chapter 8 for an introduction to file transfer operation and the use of this performance data.

The equation in Figure 40 on page 150 is used to approximate the time it takes to receive or send a file. The expansion factor E is larger than 1 when a 3X74 control unit handles also MFI traffic and/or several files concurrently. Chapter 8 shows how the Bc/B ratio is used for estimating E, and the increase in subsystem response time of MFI transactions as a result of one or more concurrent file transfer operations.

File Transfer Performance of IBM RT Personal Computer Estimated time: A + (B x E) x F seconds									
CU	WS	FDV	Size/Mode	Dir	RU/Blk	A	B	FTR	Bc/B
3174-41D	RT	Fdk	1920/CUT	Rec		1.8	0.159	6.3	0.65
				Snd		2.7	0.133	7.5	0.67

NOTES:

- A Initiation + Termination time (seconds)
  B Seconds per kbyte of file transferred
  FTR File transfer rate (kbytes per second), for more explanations, see Figure 14 in Chapter 8.
- Current RT Personal computer models are about 3x faster

Figure 40. File Transfer Performance With IBM RT Personal Computer

To enhance readability of the RT PC (ASCII) data while stored in the host files, use of the ASCII command option causes ASCII/EBCDIC translation to be performed by the file transfer program in the host. The CRLF option serves a similar purpose; it recognizes carriage return/line feed characters and transforms them into the appropriate separator characters for files in host storage.

A slow or heavily loaded host can slow down file transfers, while transfer times may also be extended by a highly fragmented near-fully loaded file device in the workstation.

APPENDIX B. ALPHAMERIC BENCHMARK DATA STREAM COMPOSITION

This appendix specifies the benchmarks used for obtaining the alphameric performance data in this guideline. When length and content of a data stream generated by an application are known, this information may enable a user to assess MFI performance trends in a specific situation.

The alphameric type A and E benchmarks are used to represent interactive transactions with the host system. They consist of an identical 40-character inbound message, and outbound data streams starting with an EW or EWA command (summarized in Figure 41 on page 153):

A-1200 Type A benchmark intended for model 2 screens: results in the display of 1200 viewable characters in 40 fields on 20 rows. Each field (alternately unprotected and protected) consists of one SBA (3 bytes), one SF (2 bytes) and 30 data characters. WCC Reset-MDT bit for EW/EWA command set 'off'. (It is ignored in EW/EWA command execution, anyway.)

A-1560 Type A for model 3 screens: causes 1560 characters to be displayed in 52 fields on 26 rows.

A-2160 Type A for model 4 screens: causes 2160 characters to be displayed in 72 fields on 36 rows.

A-4800 Type A for 3180 display stations with scroll buffer: causes 4800 characters to be displayed in 160 fields on 80 rows.

A-5760 Type A for 3290 62x160 screens: causes 5760 characters to be displayed in 192 fields on 48 rows.

E-1200 Type E benchmark intended for model 2 display stations with 7-color and extended highlighting capability. Like A-1200, its outbound data stream writes 1200 viewable characters in 40 fields on 20 rows, but each field is specified by an SFE order with three attribute pairs (8 bytes total). A display station with an extended attribute buffer (EAB) is required.

E-1560 Type E benchmark for model 3 display stations with 7-color and extended highlighting capability: displays 1560 characters in 52 fields on 26 rows; 8-byte SFE orders.

E-2160 Type E benchmark for model 4 display stations with extended highlighting capability: displays 2160 characters in 72 fields on 36 rows. SFE orders have two rather than three attribute pairs (6 bytes per order), because there are no model 4 color displays.

E-4800 Type E benchmark equivalent for A-4800 (6 bytes per SFE order).

E-5760 Type E benchmark equivalent for A-5760 (6 bytes per SFE order).

The outbound message contents of the alphameric benchmarks are ordered to correspond to a left-to-right, top-to-bottom sequence of appearance on the display screen.

The alphameric type C and C1 MFI benchmarks modify a correspondingly sized type A benchmark residing in the workstation. They have an identical 40-character inbound message and outbound data streams with a Write command. (See Figure 42 on page 154.) They emphasize the performance implications of using many PT (Program Tab) and RA (Repeat-to-Address) orders.

C-1200 Intended for use with model 2 screens: uses a WRITE command with the WCC Reset-MDT bit 'on' to modify a previous display of 1200 viewable characters created by a type A-1200 benchmark.

C-1560 Like a type C-1200 benchmark: for model 3 screens.

C-2160 Like a type C-1200 benchmark: for model 4 screens.

C-4800 Like a type C-1200 benchmark: for use in conjunction with A-4800.

C-5760 Like a type C-1200 benchmark: for use in conjunction with A-5760.

C1-1200 Just as the C-1200 benchmark, type C1-1200 modifies a display of 1200 viewable characters created by a type A-1200 benchmark. It uses a WRITE command with the Reset-MDT bit 'on', but defines about 13 percent of the fields by an SFE order with two attribute pairs (6 bytes total) for 7-color or extended highlighting. SA orders provide for changes in color or highlighting within fields. A display station with an extended attribute buffer (EAB) is required.

C1-1560 Like a type C1-1200 benchmark: for model 3 screens.

C1-2160 Like a type C1-1200 benchmark: for model 4 screens.

C1-4800 Like a type C1-1200 benchmark: for use in conjunction with A-4800.

C1-5760 Like a type C1-1200 benchmark: for use in conjunction with A-5760.

Type A and E Benchmark Data Streams									
Benchmark designation	A-1200 ##	#by	E- #by	A-1560 ##	#by	E- #by	A-2160 ##	#by	E- #by
Command: EWA/WCC	1	2	2	1	2	2	1	2	2
Orders: IC	1	1	1	1	1	1	1	1	1
SBA	40	120	120	52	156	156	72	216	216
SF - Type A	40	80	-	52	104	-	72	144	-
or SFE - Type E	40	-	320	52	-	416	72	-	432
Characters	1200	1200	1200	1560	1560	1560	2160	2160	2160
Total length, bytes		1403	1643		1823	2135		2523	2811
Result: Total rows	20			26			36		
Total fields	40 (20+20)			52 (26+26)			72 (36+36)		

Benchmark designation	A-4800 ##	#by	E- #by	A-5760 ##	#by	E- #by
Command: EWA/WCC	1	2	2	1	2	2
Orders: IC	1	1	1	1	1	1
SBA	160	480	480	192	576	576
SF - Type A	160	320	-	192	384	-
or SFE - Type E	160	-	960	192	-	1152
Characters	4800	4800	4800	5760	5760	5760
Total length, bytes		5603	6243		6723	7491
Result: Total rows	80			48		
Total fields	160 (80+80)			192 (96+96)		

NOTES:
- ## Number of commands, orders, or characters
- #by Number of bytes
- SFE 8 bytes for -1200, -1560; 6 bytes for all others

Figure 41. Types A and E Benchmark (Outbound) Contents

Type C and C1 Benchmark Data Streams

Benchmark ———>	C-1200		C1-	C-1560		C1-	C-2160		C1-
Pre-loaded with —>	(A-1200)			(A-1560)			(A-2160)		
	##	#by	#by	##	#by	#by	##	#by	#by
Command: WRITE/WCC	1	2	2	1	2	2	1	2	2
Reset-MDT bit = 'on'									
Orders: IC	1	1	1	1	1	1	1	1	1
SBA	20	60	60	26	78	78	36	108	108
SF - Type C	20	40	–	26	52	–	36	72	–
SF - Type C1	15	–	30	19	–	38	27	–	54
or SFE - Type C1	5	–	30	7	–	42	9	–	54
SA - Type C1	10	–	30	14	–	42	18	–	54
PT	10	10	10	13	13	13	18	18	18
RA	10	40	40	13	52	52	18	72	72
Characters	510	510	510	663	663	663	918	918	918
Total length, bytes		663	713		861	921		1191	1281
Result: Total rows	20			26			36		
Total fields	40 (20+20)			52 (26+26)			72 (36+36)		

Benchmark ———>	C-4800		C1-	C-5760		C1-
Pre-loaded with —>	(A-4800)			(A-5760)		
	##	#by	#by	##	#by	#by
Command: WRITE/WCC	1	2	2	1	2	2
Reset-MDT bit = 'on'						
Orders: IC	1	1	1	1	1	1
SBA	80	240	240	96	288	288
SF - Type C	80	160	–	96	192	–
SF - Type C1	60	–	120	72	–	144
or SFE - Type C1	20	–	120	24	–	144
SA - Type C1	40	–	120	48	–	144
PT	40	40	40	48	48	48
RA	40	160	160	48	192	192
Characters	2040	2040	2040	2448	2448	2448
Total length, bytes		2643	2843		3171	3411
Result: Total rows	80			144		
Total fields	160 (80+80)			192 (96+96)		

NOTES: ## Number of commands, orders, or characters
#by Number of bytes

Figure 42. Types C and C1 Benchmark (Outbound) Contents

APPENDIX C. 3X74 SELECT COMMANDS

The appropriate use of Select commands in non-SNA, channel-attached 3X74 control units is important for achieving optimum subsystem response time, and minimizing utilization of the channel and the control unit. Because the operation of the Select command in 3274 Model B and 3272 control units, and of the Select commands in non-SNA 3174 Model 1L and 3274 Model D control units differ in significant respects, they are addressed separately.

The purpose of this appendix is to help the reader use the appropriate host program options to ensure that, for a given control unit, the correct Select commands are incorporated in the data stream.

The content of this appendix is an amended version of Appendix I in Technical Bulletin **Printers Attached to 327X Control Units - Basic Performance Concepts,** G320-5906-0 (by Sumner Nash, published by the Dallas System Center, June 1982).

Select Command for 3274 Model B and 3272 Control Units

The 3274 Model B and 3272 control units process data streams in hardware logic. If the execution of a command, such as Write, RM (Read Modified), and RB (Read Buffer), requires the current content of the display buffer to be present in the control unit, some time elapses before this data is transferred from the workstation to the control unit.

Because the channel is held but not used during this time, the Select command (X'0B') was introduced to precede a Write, RM, or RB command. It will start the inbound buffer transfer, and free the channel for transfers to other control units on the channel by returning a Channel End (CE) immediately. To make sure that no commands to other devices on the same control unit are interspersed, the Write, RM, or RB command must be chained to the Select command.

Upon successful transfer of the device buffer to the control unit, a DE (Device End) is sent, signalling the host to now send the Write, RM, or RB command for immediate execution. The control unit will return CE,DE when the RM or RB operation is complete; for write operations, CE is returned after the buffer update, and DE is returned after transfer of the updated buffer to the workstation (asynchronously).

Subsequent chained commands, if any, do not need their own Select commands because the control unit buffer already contains an up-to-date copy of the workstation buffer. (Remember that an EAU command should not be chained to an EW, EWA, Write, Copy, or other EAU command.)

Select Command for 3174 and 3274 Model D Control Units

In local non-SNA 3174 and 3274 model D control units operating in CUT mode, a microprocessor processes the data stream which is partially buffered in the control unit after transfer over the channel (or before transfer for inbound data streams).

In addition to the Select RM command (X'0B', the same as for Select), there are four more commands: Select Write (X'4B'), Select RB (X'1B), Select RMP (X'2B'), and Select RBP (X'3B') (RM and RB from Position respectively).

The Select RM command initiates the preparation of an RM data stream, and causes a CE to be returned to release the channel. After scanning the buffer and building the data stream, a DE is sent to let the host know that the control unit is ready to receive the next chained command. This should be an RM command, so that the inbound transmission of the readied data stream can begin immediately, without letting the channel wait.

Any command other than RM will cause the prepared data stream to be discarded and the specified command to be executed. Although functionally the correct result is obtained, the time required to prepare the discarded RM data stream adds unnecessarily to the subsystem response time and the control unit utilization (how much depends on the display buffer content).

Correspondingly, an RB command should be chained to a Select RB command; a Write(WCC,SBA,XX,YY),RM sequence to a Select RMP; and a Write(WCC,SBA,XX,YY),RB sequence to a Select RBP. As with the Select RM command, these select commands will initiate the preparation of the correct inbound data stream. Any other commands chained to these three select commands will be rejected with CE,DE,UC,OC, Unit Check, or Operation Check.

Omission of any of the above select commands will not add to the response time but will increase channel utilization, because the channel will be held idle during preparation of the inbound data stream.

As its name implies, the 4B-Select Write command should precede a chained Write command. It will suppress preparation of an RM data stream, free the channel by returning CE, cause the keyboard to be locked, and return DE to notify the host that the command can now be sent. As soon as the Write command arrives, processing of the data stream can begin. If the Select Write command is omitted, channel utilization will be larger because the channel will be held during the time that the keyboard is being locked.

The Write Structured Field (WSF) is the only other command that may be chained to a Select Write; all others will be rejected, including the Erase Write and Erase Write Alternate commands.

Although in DFT mode data stream processing and preparation is done in the workstation rather than the control unit, there is no difference in the sequence of events as described for CUT mode. There are differences in timing, however, because the control message and data interchange between control unit and workstation differs. The do's and don'ts for CUT mode apply to DFT mode as well, only the numbers differ.

The recommendations and the performance penalties incurred by not following the recommendations are tabulated in Figure 43.

	NEXT (CHAINED) COMMAND:									
	RM		RB		Write		WSF		EW(A)	
SELECT COMMAND:	R	U	R	U	R	U	R	U	R	U
No Select command used	0	++	0	++	0	+	0	+	Rec'd	
Select RM-->	Rec'd		++	++	++	+	++	+	++	0
Select RB-->	UC,OC		Rec'd		UC,OC		UC,OC		UC,OC	
Select Write-->	UC,OC		UC,OC		Rec'd		Rec'd		UC,OC	
Select RMP-->Write*-->	Rec'd		UC,OC		UC,OC		UC,OC		UC,OC	
Select RBP-->Write*-->	UC,OC		Rec'd		UC,OC		UC,OC		UC,OC	

NOTES:

Rec'd	Recommended use
UC,OC	Unit Check, Operation Check: reject
R	Pertains to Subsystem Response Time
U	Pertains to Channel Utilization
0	"No change" in response or utilization
+	"A little increase" in response or utilization
++	"Increase" in response or utilization
-->	Indicates chained commands
*	Write command specifies buffer position only.

Figure 43. Non-SNA 3174 and 3274 Model D Select Command Sequences

Programming Support of Non-SNA 3174 and 3274 Model D Select Commands

The principal difference in Select command usage between the non-SNA 3174 model 1L (and 3274 model D), and model B and 3272 control units is the recommended use of the 0B-Select (RM). Using this command before Write and RB in model Bs, as recommended, impairs performance in 3174s and 3274 model Ds, but is not rejected.

To accommodate the modified use of the 0B-Select and the additional Select commands, changes have been made to the host programming access methods to support the control units appropriately.

This support is known as "Prepare to Read" support. Below the major operating systems are listed with an indication of available support, and how to implement it:

- **OS/VS2 (MVS) and OS/VS1**

  The SYSGEN changes are for the IODEVICE macro definition for UNIT=3277 or 3278 or 3279 or 3284 or 3286. Specify: FEATURE=PTREAD as one of the IODEVICE parameters.

  PTREAD denotes that the device is attached to a non-SNA 3174 or 3274 Model D control unit.

  A new I/O GEN must be done to include this support if it is not already in the system.

- **DOS/VSE**

  In the Automatic System Initialization IPL procedure ADD operation, place "05" as an operand for a Device Type code 3277.

  The statement will look something like: **ADD 080,3277,05**

  The "05" denotes that the display station or printer is attached to a non-SNA 3174 model 1L or 3274 model D control unit.

- **VM/SP**

  In the RIOGEN procedures, the CUTYPE under RCTLUNIT macro should be specified as follows:

Non-SNA 3174-1L:	CUTYPE = 3274
3X74-D Model:	CUTYPE = 3274
3274-B Model:	CUTYPE = 3272

APPENDIX D. LOCAL CHANNEL UTILIZATION

The data in Figure 44 can be used to estimate the average channel utilization by type A-1200 MFI benchmark transactions. The data applies to block multiplexor channels only.

For example, a 3174-1L Subsystem Control Unit in an SNA environment supporting thirty terminals, each transacting 6.7 type A-1200 transactions per minute, processes a total of 201 transactions per minute (CTR). According to the chart, this utilizes the channel about 0.9 percent, on the average.

This information is useful for estimating the total channel utilization contributed by all attached control units. For satisfactory channel performance this total should not exceed 30 percent.

	Average Channel Utilization, percent					
MFI, Data stream type	3174-1L			3274-41		
CTR —>	50	100	200	50	100	200
SNA, for A-1200	0.3	0.6	0.9	1.8	3.5	5.2
Non-SNA, for A-1200	.34	.67	1.0	.37	.74	1.1

NOTES: CTR Control Unit Transaction Rate (per minute).

For non-SNA, message lengths in excess of 14 Kb should be avoided because they cause excessive channel utilization.

Figure 44. Channel Utilization by 3X74 Control Units

ABBREVIATIONS

ACK	The acknowledge character (BSC)
AEA	ASCII emulation adapter
AID	Attention identifier
APA	All points addressable
APL	A programming language
ASCII	American National Standard Code for Information Interchange
BMS	Basic mapping support (CICS)
BSC	Binary synchronous communications
CCITT	International Telegraph and Telephone Consultative Committee
CICS	Customer information control system
CP	Control program
CR	The 'carriage return' character (3270 & SCS)
CTR	Control unit transaction rate
CU	Control unit
CUT	Control unit terminal mode
DCA	Device cluster adapter (3274)
DFT	Distributed function terminal mode
DLE	The data link escape character (BSC)
DS	Data stream
DSC	3270 Data stream compatibility
EAB	Extended attribute buffer
EAU	Erase all unprotected (3270 command)
EDGAR	ACMS editor
EDS	Extended data stream
EM	The 'end of medium' character (3270 & SCS)
ENQ	The enquiry character (BSC)

EOT	The 'end of transmission' character (BSC)
EPC	Early Print Complete
ETB	The 'end of transmission block' character (BSC)
ETX	The 'end of text' character (BSC)
EUA	Erase unprotected to address (3270 order)
EW	Erase write (3270 command)
EWA	Erase write alternate (3270 command)
FDV	File device type
FF	The 'form feed' character (3270 & SCS)
FT	File transfer
FTR	File transfer rate
GDDM	Graphical data display manager
IAT	Inter arrival time
IC	Insert cursor (3270 order)
IMS	Information management system
Kb	Kilobyte; 1024 bytes
kbs	kilo (1000) bits per second, as in 19.2 kbs
kbytes	kilo (1000) bytes
LPI	Lines per inch (printers)
LPM	Lines per minute (print speed)
LPS	Load programmed symbols (3270 structured field)
Mb	Megabyte; 1,048,576 bytes
MDT	Modified data tag (3270)
MF	Modify field (3270 order)
MFI	Main frame interactive
NCP	Network control program
NDCA	New device cluster adapter (3174)
NL	The 'newline' character (3270 & SCS)

PC	IBM Personal Computer
ppi	Picture elements (pels) per inch
PS	Programmed symbols (3270)
PT	Program tab (3270 order)
PU	Physical unit (SNA)
RA	Repeat to address (3270 order)
RAM	Random access memory (main storage)
RB	Read buffer (3270 command)
RM	Read modified (3270 command)
RT	Response time (to last character)
RTR	Ring transaction rate (in Token-Ring Network)
SA	Set attribute (3270 order)
SBA	Set buffer address (3270 order)
SCP	System control program
SCS	SNA character string
SDLC	Synchronous data link control
SF	Start field (3270 order)
SFE	Start field extended (3270 order)
SIO	Start input/output operation
SNA	Systems network architecture
SPC	Set Printer Characteristics (structured field)
SPF	System Productivity Facility, or Structured Programming Facility
SS	Subsystem
TCA	Terminal communications adapter
TMA	Terminal multiplexer adapter (3174)
TP	Teleprocessing
TRN	Token-Ring Network
TSO	Time sharing option

VM Virtual machine

WACK The 'wait before transmit' positive acknowledgement (BSC)

WCC Write control character (3270)

WSF Write structured field (3270 command)

XEDIT A CMS editor

Performance Guidelines
for 3X74-- Attached Workstations

Publication No. ZZ20-4167-7

READER'S COMMENT FORM

This manual is part of a library that serves as a reference source for systems analysts, programmers, and operators of IBM systems. You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Note: Copies of IBM publications are not stocked at the location to which this form is addressed. Please direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

If you wish a reply, give your name, company, mailing address, and date:

____

____

____

____

What is your occupation? ______________________________

Number of latest Newsletter associated with this publication: ______________________________

Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A. (Elsewhere, an IBM office or representative will be happy to forward your comments or you may mail directly to the address in the Edition Notice on the back of the title page.)

Reader's Comment Form

Fold and tape | **Please Do Not Staple** | Fold and tape

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Dept. G37 Building 002
P.O. Box 12195
Research Triangle Park, N.C. 27709-9990

Fold and tape | **Please Do Not Staple** | Fold and tape

IBM®

ZZ20-4167-7

PRINTED IN U.S.A. ZZ20-4167-7

ZZ20-4167-07

IBM®